

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,)	
)	
Plaintiffs,)	
)	
v.)	C.A. No.
)	
LANCIUM LLC, MICHAEL T.)	JURY TRIAL DEMANDED
MCNAMARA, and RAYMOND E. CLINE,)	
JR.)	
)	
Defendants.)	

COMPLAINT

Plaintiffs BearBox LLC (“BearBox”) and Austin Storms (collectively, “Plaintiffs”) bring this action against Lancium LLC (“Lancium”), Michael T. McNamara, and Raymond E. Cline, Jr. (collectively “Defendants”) to correct the inventorship of U.S. Patent No. 10,608,433 (the “’433 Patent”) and to recover damages, injunctive relief, declaratory relief, and other remedies for Defendants’ wrongful actions to obtain, misuse, disclose, and claim as their own Plaintiffs’ proprietary cryptocurrency mining technology. Plaintiffs further allege as follows:

INTRODUCTION

1. This case is about the Defendants’ theft of inventions that rightfully belong to Plaintiffs.
2. Plaintiffs developed an energy-efficient cryptocurrency mining system and related methods that reduce the inefficiency and environmental impact of energy-expensive mining operations by better utilizing available energy resources to increase stability of the energy grid, minimize a mining operation’s impact on peak-demand, and also alleviate electricity

undersupply and/or oversupply conditions (the “BearBox Technology”). The BearBox Technology can be used to mine cryptocurrency, such as Bitcoin.

3. The Defendants induced the Plaintiffs to disclose the BearBox Technology to them under the guise of a possible business deal between Defendants and Plaintiffs to jointly commercialize the BearBox Technology. Before disclosing the BearBox Technology to Defendants, Plaintiffs obtained assurances of confidentiality from Defendants.

4. Rather than keeping the BearBox Technology confidential and using it in furtherance of a business relationship with Plaintiffs, the Defendants stole the BearBox Technology from Plaintiffs by converting and misappropriating it and claiming it as their own. Defendants filed a U.S. patent application that wrongfully disclosed the BearBox Technology to the U.S. Patent and Trademark Office and ultimately to the public. And by obtaining the ’433 Patent with claims directed to the BearBox Technology, the Defendants have wrongfully obtained a patent covering the BearBox Technology and wrongfully claimed the BearBox Technology as their own.

5. Plaintiffs bring this action to correct the named inventors on the ’433 Patent. The inventions claimed in the ’433 Patent are inventions conceived by Storms, founder and president of BearBox.

PARTIES

6. Plaintiff BearBox LLC (“BearBox”) is a limited liability company organized and existing under the laws of Louisiana with its principal place of business at 4422 Highway 22, Mandeville, Louisiana 70471.

7. Plaintiff Austin Storms is an individual residing in Mandeville, Louisiana.

8. On information and belief, Defendant Lancium is a Delaware limited liability company with its principal place of business at 6006 Thomas Rd, Houston, Texas 77041. On information and belief, Lancium has a registered agent capable of accepting service in this district, Harvard Business Services, Inc. with a place of business at 16192 Coastal Highway, Lewes, DE 19958.

9. On information and belief, Defendant Michael T. McNamara is the Chief Executive Officer and a founder of Lancium and resides in Newport Beach, California. Defendant McNamara is named as a purported inventor on the face of the '433 Patent.

10. On information and belief, Defendant Raymond E. Cline, Jr. is the Chief Computing Officer of Lancium and resides in Houston, Texas. Defendant Cline is named as a purported inventor on the face of the '433 Patent.

JURISDICTION

11. This is an action seeking correction of the named inventors of a United States patent under 35 U.S.C. § 256. As such, this action arises under the laws of the United States.

12. This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under an Act of Congress relating to patents, specifically 35 U.S.C. § 256.

13. This Court further has subject matter jurisdiction under 28 U.S.C. §§ 1331 because the matter arises under the trade secret laws of the United States, 18 U.S.C. § 1836.

14. The Court has supplemental jurisdiction under 28 U.S.C. § 1367 over all asserted claims under state law because those claims are so related to the claims in this action that arise under federal law that they form part of the same case or controversy.

15. The Court also has jurisdiction pursuant to 28 U.S.C. § 1332, as complete diversity of citizenship exists among the parties, and the amount in controversy exceeds \$75,000. Plaintiff BearBox is a citizen of the State of Louisiana because it is organized under the laws of the State of Louisiana and has its principal place of business in the State of Louisiana. Plaintiff Storms is a citizen of the State of Louisiana because he resides in the State of Louisiana. In contrast, none of the Defendants are citizens of the State of Louisiana. Defendant Lancium is a citizen of the States of Delaware and Texas because it is organized under the laws of the State of Delaware and has its principal place of business in the State of Texas. Defendant McNamara is a citizen of the State of California because he resides in the State of California. Defendant Cline is a citizen of the State of Texas because he resides in the State of Texas. Therefore, because the Plaintiffs are both citizens of the State of Louisiana (and no other states) for purposes of diversity jurisdiction, and none of the Defendants are citizens of the State of Louisiana, complete diversity exists among the parties.

16. This Court has general personal jurisdiction over Lancium because it is organized under the laws of the State of Delaware and because it maintains an ongoing presence in this District at least through its registered agent.

17. This Court has specific personal jurisdiction over each of Defendants McNamara and Cline at least under Title 6 of the Delaware Code, § 18-109(a).

18. On information and belief, Defendant McNamara is the Chief Executive Officer of Lancium. On information and belief, as the Chief Executive Offer, McNamara participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

19. McNamara is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from the interests of Lancium and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claims against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant McNamara together. Plaintiffs' claims against Defendant McNamara arise out of his exercise of his powers as Chief Executive Officer of Lancium.

20. On information and belief, Defendant Cline is the Chief Computing Officer of Lancium. On information and belief, as the Chief Computing Officer, Cline participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

21. Cline is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from Lancium's interest and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claim against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant Cline together. Plaintiffs' claims against Defendant Cline arise out of his exercise of his powers as Chief Computing Officer of Lancium.

22. The actions of Defendants McNamara and Cline establish sufficient minimum contacts with Delaware under Delaware law and the United States Constitution to give this Court personal jurisdiction over each of them.

23. As described below, each Defendant has committed acts giving rise to this action.

VENUE

24. Venue is proper in this District under 28 U.S.C. § 1391(b)(3) because there is no district in which an action may otherwise be brought as provided in § 1391(b) and Defendant Lancium is subject to the Court’s personal jurisdiction with respect to this action.

PLAINTIFFS’ PROPRIETARY CRYPTOCURRENCY MINING TECHNOLOGY

25. As of 2018, the amount of energy required to process computer algorithms to mine cryptocurrencies like Bitcoin was three times greater than the energy required to physically mine gold. Conventional mining of “copper, gold, platinum, and rare earth oxides are 4, 5, 7, and 9 megajoules to generate one U.S. dollars,” while “it costs an average of 17 megajoules to mine \$1 worth of bitcoin.”¹ The large amount of energy required to mine cryptocurrencies can make such mining financially prohibitive, and even when financially lucrative, the large energy requirements make cryptocurrency mining harmful to the global environment, with studies showing carbon dioxide emissions from cryptocurrency mining “single-handedly rais[ing] global temperatures by 2 degrees by 2023.” *Id.*

26. At the same time, some forms of electrical power generation are terribly inefficient. When producers of electrical power are unable to quickly adjust their operations in response to dynamically changing grid conditions, these producers frequently sell power at low or even negative prices until demand and market prices increase.

27. Because cryptocurrency mining is a computationally demanding process, it requires significant energy. As a result, industrial-scale cryptocurrency mining places a large

¹ <https://www.marketwatch.com/story/mining-bitcoin-is-3-times-more-expensive-than-mining-gold-research-paper-finds-2018-11-06>

energy burden on the power grid, driving demand and costs as well as increasing the likelihood of grid component failure.

28. In late 2018 and early 2019, Austin Storms sought to address these problems by developing energy-efficient cryptocurrency mining systems and methods that reduce the environmental impact of energy-intensive mining operations. Storms conceived of a system that better uses available energy resources to increase the stability of the energy grid, minimize a mining operation's impact on peak-demand, and alleviate electricity undersupply and/or oversupply conditions, all while decreasing the overall energy costs of the mining operation and increasing its profitability.

29. Austin Storms conceived of and developed the BearBox Technology. Storms is the president and founder of BearBox. The BearBox Technology includes hardware and software components. Structurally, the BearBox Technology includes a housing for a plurality of miners (such as ASICs, graphics cards, or the like) under the direction of a smart controller(s).

30. The smart controller monitors various external factors, such as current and expected energy demand and pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the system may determine whether conditions are appropriate to mine cryptocurrency and, if so, subsequently mines the cryptocurrency. Optionally, the system also includes other components for cooling, air-filtration, and related features.

31. In the BearBox Technology, a controller (such as a power distribution unit, network interface, or the like) monitors various external factors, such as current and expected energy demand/pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the controller(s) determines appropriate times to mine

cryptocurrency in accordance with a desired performance strategy (for example, profitability thresholds). At the appropriate times, the controller initiates mining, for example, by powering on the miners.

DEFENDANTS WRONGFULLY CLAIM THE BEARBOX TECHNOLOGY AS THEIR OWN

32. In May 2019, Storms attended the Fidelity FCAT Mining Summit in Boston, Massachusetts on behalf of BearBox to promote the BearBox Technology and seek potential customers for his revolutionary system.

33. While at the conference, Storms met Defendant McNamara. Defendant McNamara showed immediate interest in the BearBox Technology. Under the rouse of a potential business relationship, McNamara pumped Storms for details about the BearBox Technology over the course of several exchanges, which included conversations, emails, and text messages about the BearBox Technology. Storms took McNamara to dinner where McNamara continued to pump Storms for details about the BearBox Technology. At all times before and during Storms's disclosure of this information, Storms told McNamara that the BearBox Technology was confidential, and Storms relied on McNamara's good faith assurances that he would keep confidential the information he received from Storms about the BearBox Technology.

34. Following the conference, McNamara continued to press Storms for additional details about the BearBox Technology via text messaging and email. Again relying on Defendant McNamara's assurances of confidentiality, Storms provided annotated system diagrams, component specifications, and modeled data sets to mimic real-world Bitcoin and energy prices. Storms included express confidentiality notices in his communications with Defendant McNamara.

35. After Storms disclosed the BearBox Technology to McNamara, McNamara abruptly ended all communications with Storms.

36. Storms last communicated with McNamara on May 9, 2019 via e-mail, and after sending that message, Storms did not hear from McNamara again.

37. At that time, Storms understood that McNamara was not interested in investing in the BearBox Technology. He had no reason to suspect that McNamara would steal the BearBox Technology and claim it as his own.

38. On information and belief, Defendants filed U.S. provisional patent application No. 62/927,119 on October 28, 2019, naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

39. In addition to falsely claiming to be the inventors of the inventions disclosed in the application, Defendants wrongfully disclosed, without authorization, the confidential BearBox Technology to the United States Patent and Trademark Office.

40. Likewise, on December 4, 2019, Defendants filed U.S. Patent Application Serial No. 16/702,931, once again naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

41. The '433 Patent issued on March 31, 2020 naming Defendants McNamara and Cline as the sole purported inventors on the face of the patent. A true and correct copy of the '433 Patent is attached hereto as Exhibit A.

42. The inventions claimed in the '433 patent encompass the BearBox Technology, yet Defendants falsely identified themselves as the inventors of the claimed inventions, when, in fact, Storms is the sole inventor of the claimed inventions.

43. On information and belief, McNamara and Cline assigned their purported rights in the '433 patent to Lancium. On information and belief, at all times, Lancium was aware that McNamara and Cline, both officers of Lancium, were not the rightful inventors of the BearBox Technology disclosed in the patent and the inventions claimed in the patent.

44. Defendants McNamara and Cline each submitted signed declarations falsely swearing that they were “an original joint inventor” of the claimed subject matter. A true and correct copy of Defendant McNamara’s and Defendant Cline’s declarations are attached as Exhibit B.

45. On August 14, 2020, Lancium filed a lawsuit in the U.S. District Court for the Western District of Texas against Layer1 Technologies, Inc. (“Layer1”) asserting that Layer1 infringes the '433 patent. That case is captioned *Lancium LLC v. Layer1 Technologies, Inc.*, Case No. 6:20-cv-739 (W.D. Texas) (the “Layer1 Lawsuit”).

46. As part of the Layer1 Lawsuit, Defendants falsely asserted that McNamara and Cline are the sole inventors of the inventions claimed in the '433 patent.

47. Plaintiffs became aware of Defendants’ wrongful use of the BearBox Technology on or about August 17, 2020, when they learned about the Layer1 Lawsuit through a press release dated August 14, 2020, posted by Lancium on PRNewswire. That press release is available at the following URL: <https://www.prnewswire.com/news-releases/controllable-load-resource-clr-market-leader-lancium-files-patent-infringement-lawsuit-against-layer1-301112687.html>.

48. Before seeing the August 14, 2020 press release, Plaintiffs were unaware of Defendants’ wrongful use of the BearBox Technology and was unaware of the '433 patent.

49. On March 5, 2021, Lancium and Layer1 entered a Stipulation to Dismiss with Prejudice in the Layer1 Lawsuit. According to the stipulation, the parties had entered a Settlement Agreement to resolve the Layer1 Lawsuit.

50. According to a press release issued by Lancium on March 8, 2021, Lancium and Layer 1 “have entered into a mutually beneficial partnership. Layer1 has licensed Lancium’s intellectual property and Lancium will provide Smart Response™ software and services to Layer1.” The press release is available at the following URL: <https://www.prnewswire.com/news-releases/lancium-and-layer1-settle-patent-infringement-suit-301242602.html>

51. On information and belief, as part of the Settlement Agreement between Lancium and Layer1 to settle the Layer1 Lawsuit, Lancium received and continues to receive valuable consideration from Layer1, all of which rightly belongs to Plaintiffs, the rightful owners of the inventions claimed in the ’433 Patent.

COUNT I
CORRECTION OF INVENTORSHIP FOR THE ’433 PATENT:
AUSTIN STORMS AS SOLE INVENTOR

52. Plaintiffs incorporate the above paragraphs by reference.

53. Storms is the sole inventor of the subject matter claimed in the ’433 Patent.

54. Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the ’433 patent and the currently listed inventors on the ’433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms or BearBox.

55. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the ’433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT II
IN THE ALTERNATIVE, CORRECTION OF INVENTORSHIP FOR THE '433
PATENT: AUSTIN STORMS AS JOINT INVENTOR WITH THE CURRENTLY
NAMED INVENTORS

56. Plaintiffs incorporates the above paragraphs by reference.

57. In the alternative, Storms is a joint inventor of the subject matter claimed in the '433 Patent and should be added to the individuals currently named as inventors on the '433 Patent.

58. Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the '433 patent and the currently listed inventors on the '433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms.

59. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the '433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT III
CONVERSION BY LANCIUM, MCNAMARA, AND CLINE

60. Plaintiffs incorporates the above paragraphs by reference.

61. Austin Storms, in his capacity as founder and President of BearBox, conceived, developed, and reduced to practice the BearBox Technology. Plaintiffs own the BearBox Technology, related know-how, and related intellectual property. Plaintiffs owned this property during all relevant time periods in this suit. Information on the BearBox Technology was provided to Defendants solely for the purposes of evaluation for a potential business relationship and under strict confidentiality obligations.

62. Defendants assumed dominion and control over the BearBox Technology by claiming it as their own in the '433 patent. Through their wrongful conduct in obtaining the '433 Patent and claiming the BearBox Technology as their own, the Defendants have wrongfully obtained the purported ability to exclude Plaintiffs and others from using the BearBox Technology. This constitutes unauthorized and unlawful conversion by Defendants.

63. As a result of Defendants' wrongful actions, Plaintiffs will suffer imminent and irreparable damages in an amount to be proven at trial. In particular, Plaintiffs have been damaged by losing valuable intellectual property from which Plaintiffs would have derived substantial revenue via licensing and/or selling patented products.

**COUNT IV
UNJUST ENRICHMENT BY LANCIUM, MCNAMARA, AND CLINE**

64. Plaintiffs incorporate the above paragraphs by reference.

65. Plaintiffs conferred a benefit on Defendants by providing them valuable intellectual property about cryptocurrency mining systems and related confidential information and materials under the boundaries of a potential collaboration between BearBox and Lancium.

66. Defendants accepted that cryptocurrency mining intellectual property and, indeed, continuously asked Storms to provide more information and materials, having recognized the benefit that Defendants received by having access to the BearBox Technology.

67. Defendants accepted and retained the BearBox Technology, and used it to their own advantage, at Plaintiffs' expense.

68. Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct. In particular, Defendants have unlawfully used Plaintiffs' property by asserting inventorship over the BearBox Technology, and deriving an unjust benefit from

exploiting Storms's cryptocurrency mining inventions. It would be inequitable for Defendants to retain these benefits under these circumstances.

69. Plaintiffs have incurred, and continue to incur, detriment in the form of loss of money and property as a result of Defendants' wrongful use of Plaintiffs' intellectual property, including the right to any patent based on their own intellectual property and any royalties derived from that intellectual property. The intellectual property, including the right to any patents based on Plaintiffs' intellectual property and to any patent documents (including assignment documents), U.S. and foreign, are unique and there is no adequate remedy at law.

70. The harm to Plaintiffs is continuous, substantial, and irreparable.

COUNT V
TRADE SECRET MISAPPROPRIATION UNDER
FEDERAL DEFEND TRADE SECRETS ACT, 18 U.S.C. § 1836

71. Plaintiffs incorporate the above paragraphs by reference.

72. At all relevant times, Plaintiffs and Defendants were engaged in interstate commerce. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods in interstate commerce throughout the United States.

73. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about the BearBox Technology relating to cryptocurrency mining systems and related methods from Plaintiff. The confidential information included detailed technical information about the BearBox Technology. Such information constitutes trade secrets of substantial economic value. This confidential technical information was not known to the public or to other persons who can obtain economic value from its

disclosure or use. This information constituted a trade secret under the Defend Trade Secrets Act, 18 U.S.C. §§ 1836.

74. Defendants knew the technical information they received from Austin Storms about the BearBox Technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, for internal use only, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship. Plaintiffs took reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its offices and computers.

75. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent constituted a misappropriation of Plaintiffs' trade secrets. Defendants acquired the confidential information under circumstances giving rise to a duty to maintain the secrecy of the trade secret or limit the use of the trade secret.

76. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent was an unauthorized disclosure in breach of Defendants' agreement to maintain the secrecy of this information or otherwise limit its use. As recently as December 4, 2019, Defendants have wrongly sworn that Plaintiffs' confidential technical information is their own by declaring to the United States Patent and Trademark Office that they were the original and joint inventors of the inventions claimed in the '433 patent. Even ignoring that Defendants have improperly claimed this information as their own, through their public filings Defendants have disclosed the confidential technical information without the express or implied consent of

Plaintiffs. The acts of taking Plaintiffs' confidential information, claiming it as Defendants' own, and submitting it in public filings were improper means in breach of a confidential relationship. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

77. As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, damages for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability for a reasonable royalty for the Defendants' unauthorized disclosure or use of the trade secrets.

78. The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D).

**COUNT VI
TRADE SECRET MISAPPROPRIATION UNDER
TEXAS CIVIL PRACTICE AND REMEDIES CODE § 134A**

79. Plaintiffs incorporate the above paragraphs by reference.

80. At all relevant times, Plaintiffs and Defendants were engaged in trade or commerce within the meaning of Texas Civil Practice and Remedies c. § 134A, the Texas Uniform Trade Secrets Act. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods.

81. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about the BearBox Technology relating to cryptocurrency mining systems and related methods from Plaintiff. The confidential information included detailed technical information about the BearBox Technology. Such information constitutes trade secrets of substantial economic value. This confidential technical information was not known to the public or to other persons who can obtain economic value from its disclosure or use. This information constituted a trade secret.

82. Defendants knew the technical information they received from Austin Storms about the BearBox Technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, for internal use only, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship. Plaintiffs took reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its offices and computers.

83. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent constituted a misappropriation of Plaintiffs' trade secrets. Defendants acquired the confidential information under circumstances giving rise to a duty to maintain the secrecy of the trade secret or limit the use of the trade secret.

84. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent was an unauthorized disclosure in breach of Defendants' agreement to maintain the secrecy of this information or otherwise limit its use. As recently as December 4, 2019,

Defendants have wrongly sworn that Plaintiffs' confidential technical information is their own by declaring to the United States Patent and Trademark Office that they were the original and joint inventors of the inventions claimed in the '433 patent. Even ignoring that Defendants have improperly claimed this information as their own, through their public filings Defendants have disclosed the confidential technical information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, claiming it as Defendants' own, and submitting it in public filings were improper means in breach of a confidential relationship. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

85. As a result of the misappropriation of Plaintiffs' trade secrets and intended immediate use by the Defendants, Plaintiffs' have suffered and will suffer damages in an amount to be proven at trial.

COUNT VII NEGLIGENT MISREPRESENTATION BY LANCIMUM AND MCNAMARA

86. Plaintiffs incorporate the above paragraphs by reference.

87. In connection with the potential work involving cryptocurrency mining systems and related methods, Storms told Defendant McNamara that the cryptocurrency mining systems and related methods were proprietary to Plaintiffs and not to be used or shared outside of Lancium. Defendant McNamara gave his word that he would abide by this confidentiality. On information and belief, Defendant McNamara agreed to keep the BearBox Technology confidential despite later recklessly incorporating the BearBox Technology into his own patent applications and swearing, as recently as December 4, 2019, that he is an inventor of the

BearBox Technology. Storms relied on Defendant McNamara's assurances of confidentiality and continued to share details about the BearBox Technology with Defendants.

88. If Plaintiffs had known that Defendants would secretly incorporate the BearBox Technology into Defendants' own patent applications to claim them as Defendants' intellectual property, Plaintiffs would not have continued working with and sharing intellectual property with Defendants.

89. Plaintiffs suffered a pecuniary loss based on this reliance including the loss of potential patent rights, and the costs of Plaintiffs' know-how converted under the guise of a potential business relationship.

JURY DEMAND

90. Under Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiffs respectfully demand a trial by jury on all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, BearBox respectfully requests the following relief:

A. An order that the Director of the United States Patent and Trademark Office correct the inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

B. Alternatively, an order that Defendants sign the requisite documents to correct inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

C. A declaration that Austin Storms is the sole inventor, or, in the alternative, is a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

D. A preliminary and a permanent injunction enjoining Defendants Lancium, McNamara, and Cline from asserting that McNamara or Cline are inventors of the '433 Patent in violation of the United States federal patent laws;

E. An order that Defendants immediately transfer to Plaintiffs all right, title, and interest in all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

F. An order for a constructive trust over all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

G. Financial relief including damages, consequential damages, disgorgement of Defendants' ill-gotten profits, Defendants' unjust enrichment, restitution, reasonable royalty damages, lost profits damages, reliance damages, and/or all other appropriate financial relief, all in an amount to be determined at trial, with interest;

H. An award of the amount by which Defendants have been unjustly enriched by their actions set forth in this Complaint and their purported ownership of patents covering Plaintiffs' intellectual property;

I. An award of exemplary, enhanced, and/or punitive damages as permitted by law, including under 18 U.S.C. § 1836(b)(3)(C) and Texas Business and Commerce Code § 17E;

J. A finding that this is an exceptional case warranting imposition of attorney fees

K. An award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(D); and

L. An award of such further relief at law or in equity, such as preliminary and/or

ASHBY & GEDDES

/s/ Andrew C. Mayo

Andrew C. Mayo (#5207)
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19899
(302) 654-1888
amayo@ashbygeddes.com

*Attorneys for Plaintiffs
BearBox LLC and Austin Storms*

Of Counsel:

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati, III
MARSHALL, GERSTEIN & BORUN LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
(312) 474-6300

Dated: April 14, 2021

EXHIBIT A

(12) **United States Patent**
McNamara et al.

(10) **Patent No.:** **US 10,608,433 B1**
 (45) **Date of Patent:** **Mar. 31, 2020**

(54) **METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT**

7,143,300 B2 11/2006 Potter et al.
 7,647,516 B2 1/2010 Ranganathan et al.
 (Continued)

FOREIGN PATENT DOCUMENTS

(71) Applicant: **Lancium LLC**, Houston, TX (US)

CN 103163904 A 6/2013
 KR 20090012523 A 2/2009
 WO 2015199629 A1 12/2015

(72) Inventors: **Michael T. McNamara**, Newport Beach, CA (US); **Raymond E. Cline, Jr.**, Houston, TX (US)

OTHER PUBLICATIONS

(73) Assignee: **Lancium LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Bird et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," National Renewable Energy Lab (NREL), Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.

(Continued)

(21) Appl. No.: **16/702,931**

Primary Examiner — Christopher E. Everett

(22) Filed: **Dec. 4, 2019**

(74) *Attorney, Agent, or Firm* — McDonnell Boehnen Hulbert & Berghoff LLP

Related U.S. Application Data

(60) Provisional application No. 62/927,119, filed on Oct. 28, 2019.

(57) **ABSTRACT**

(51) **Int. Cl.**
H02J 3/14 (2006.01)
H02J 3/00 (2006.01)
G06F 1/3203 (2019.01)

Examples relate to adjusting load power consumption based on a power option agreement. A computing system may receive power option data that is based on a power option agreement and specify minimum power thresholds associated with time intervals. The computing system may determine a performance strategy for a load (e.g., set of computing systems) based on a combination of the power option data and one or more monitored conditions. The performance strategy may specify a power consumption target for the load for each time interval such that each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The computing system may provide instructions the set of computing systems to perform one or more computational operations based on the performance strategy.

(52) **U.S. Cl.**
 CPC **H02J 3/14** (2013.01); **G06F 1/3203** (2013.01); **H02J 3/008** (2013.01)

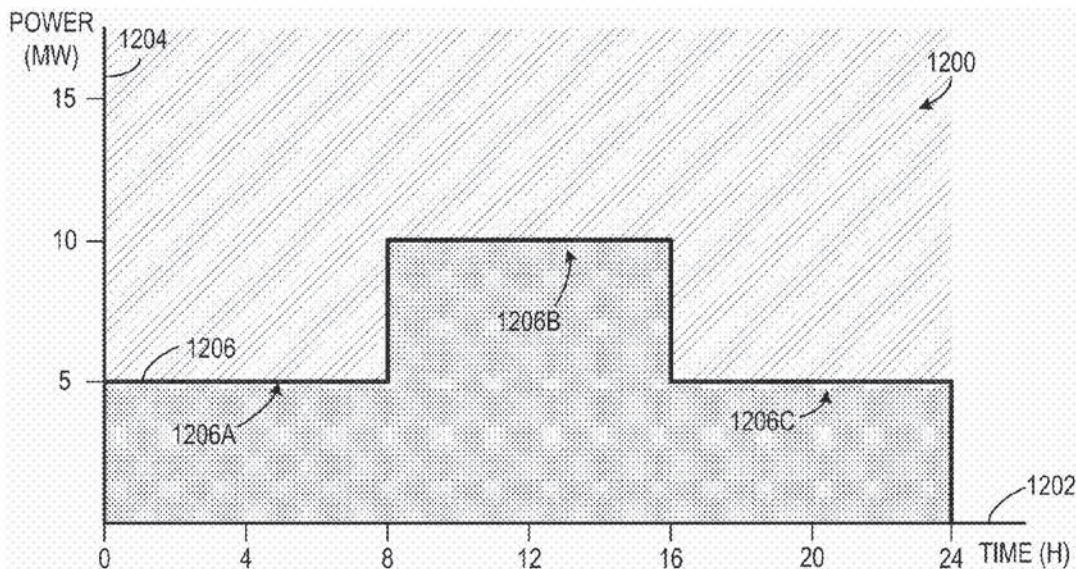
(58) **Field of Classification Search**
 CPC H02J 3/14; H02J 3/008; G06F 1/3203
 See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,288,456 B1 9/2001 Cratty
 6,633,823 B2 10/2003 Bartone et al.

20 Claims, 16 Drawing Sheets



US 10,608,433 B1

Page 2

(56)

References Cited

U.S. PATENT DOCUMENTS

7,702,931	B2	4/2010	Goodrum et al.	
7,779,276	B2	8/2010	Bolan et al.	
7,861,102	B1	12/2010	Ranganathan et al.	
7,921,315	B2	4/2011	Langgood et al.	
7,970,561	B2	6/2011	Pfeiffer	
8,001,403	B2	8/2011	Hamilton et al.	
8,006,108	B2	8/2011	Brey et al.	
8,214,843	B2	7/2012	Boss et al.	
8,374,928	B2	2/2013	Gopisetty et al.	
8,447,993	B2	5/2013	Greene et al.	
8,571,820	B2	10/2013	Pfeiffer	
8,627,123	B2	1/2014	Jain et al.	
8,700,929	B1*	4/2014	Weber	G06F 30/13 713/310
8,789,061	B2	7/2014	Pavel et al.	
8,799,690	B2	8/2014	Dawson et al.	
9,003,211	B2	4/2015	Pfeiffer	
9,003,216	B2	4/2015	Sankar et al.	
9,026,814	B2	5/2015	Aasheim et al.	
9,207,993	B2	12/2015	Jain	
9,218,035	B2	12/2015	Li et al.	
9,552,234	B2	1/2017	Boldyrev et al.	
10,367,353	B1	7/2019	McNamara et al.	
10,367,535	B2	7/2019	Corse et al.	
10,444,818	B1	10/2019	McNamara et al.	
10,452,127	B1	10/2019	McNamara et al.	
10,497,072	B2	12/2019	Hooshmand et al.	
2002/0072868	A1	6/2002	Bartone et al.	
2003/0074464	A1	4/2003	Bohrer et al.	
2005/0203761	A1*	9/2005	Barr	G06F 1/26 713/320
2006/0161765	A1	7/2006	Cromer et al.	
2008/0030078	A1	2/2008	Whitted et al.	
2008/0094797	A1	4/2008	Coglitore et al.	
2009/0055665	A1	2/2009	Maglione et al.	
2009/0070611	A1*	3/2009	Bower, III	G06F 1/3203 713/322
2009/0089595	A1*	4/2009	Brey	G06F 1/3203 713/300
2010/0211810	A1	8/2010	Zacho	
2010/0328849	A1	12/2010	Ewing et al.	
2011/0072289	A1	3/2011	Kato	
2012/0000121	A1	1/2012	Swann	
2012/0072745	A1	3/2012	Ahluwalia et al.	
2012/0300524	A1	11/2012	Fornage et al.	
2012/0324259	A1	12/2012	Aasheim et al.	
2013/0006401	A1	1/2013	Shan	
2013/0063991	A1	3/2013	Xiao et al.	
2013/0086404	A1*	4/2013	Sankar	G06F 1/305 713/324
2013/0187464	A1	7/2013	Smith et al.	
2013/0227139	A1	8/2013	Suffling	
2013/0306276	A1	11/2013	Duchesneau	
2014/0137468	A1	5/2014	Ching	
2014/0379156	A1	12/2014	Kamel et al.	
2015/0121113	A1	4/2015	Ramamurthy et al.	
2015/0155712	A1	6/2015	Mondal	
2015/0229227	A1	8/2015	Aeloiza et al.	
2015/0277410	A1	10/2015	Gupta et al.	
2015/0278968	A1	10/2015	Steven et al.	
2016/0170469	A1	6/2016	Schgal et al.	
2016/0172900	A1	6/2016	Welch, Jr.	
2016/0187906	A1*	6/2016	Bodas	G05B 15/02 700/287

2016/0198656	A1	7/2016	McNamara et al.
2016/0212954	A1	7/2016	Argento
2016/0324077	A1	11/2016	Frantzen et al.
2017/0023969	A1	1/2017	Shows et al.
2017/0104336	A1	4/2017	Elbsat et al.
2017/0261949	A1	9/2017	Hoffmann et al.
2018/0144414	A1	5/2018	Lee et al.
2018/0202825	A1	7/2018	You et al.
2018/0240112	A1	8/2018	Castinado et al.
2018/0366978	A1	12/2018	Matan et al.
2018/0367320	A1	12/2018	Montalvo
2019/0052094	A1	2/2019	Pmsvsv et al.
2019/0168630	A1	6/2019	Mrlik et al.
2019/0258307	A1	8/2019	Shaikh et al.
2019/0280521	A1	9/2019	Lundstrom et al.
2019/0318327	A1	10/2019	Sowell et al.
2019/0324820	A1	10/2019	Krishnan et al.

OTHER PUBLICATIONS

EPEX Spot, "How They Occur, What They Mean," https://www.epexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.

Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.

Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017955, dated Apr. 30, 2018, 22 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017950, dated May 31, 2018, 15 pages.

Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.

Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.

Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.

Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.

Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.

Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.

Rahimi, Farrokh, "Using a Transactive Energy Framework," IEEE Electrification Magazine, Dec. 2016, pp. 23-29.

Soluna, "Powering the Block Chain," Aug. 2018, version 1.1, 29 pages.

Wilson, Joseph Nathanael, "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016, 154 pages.

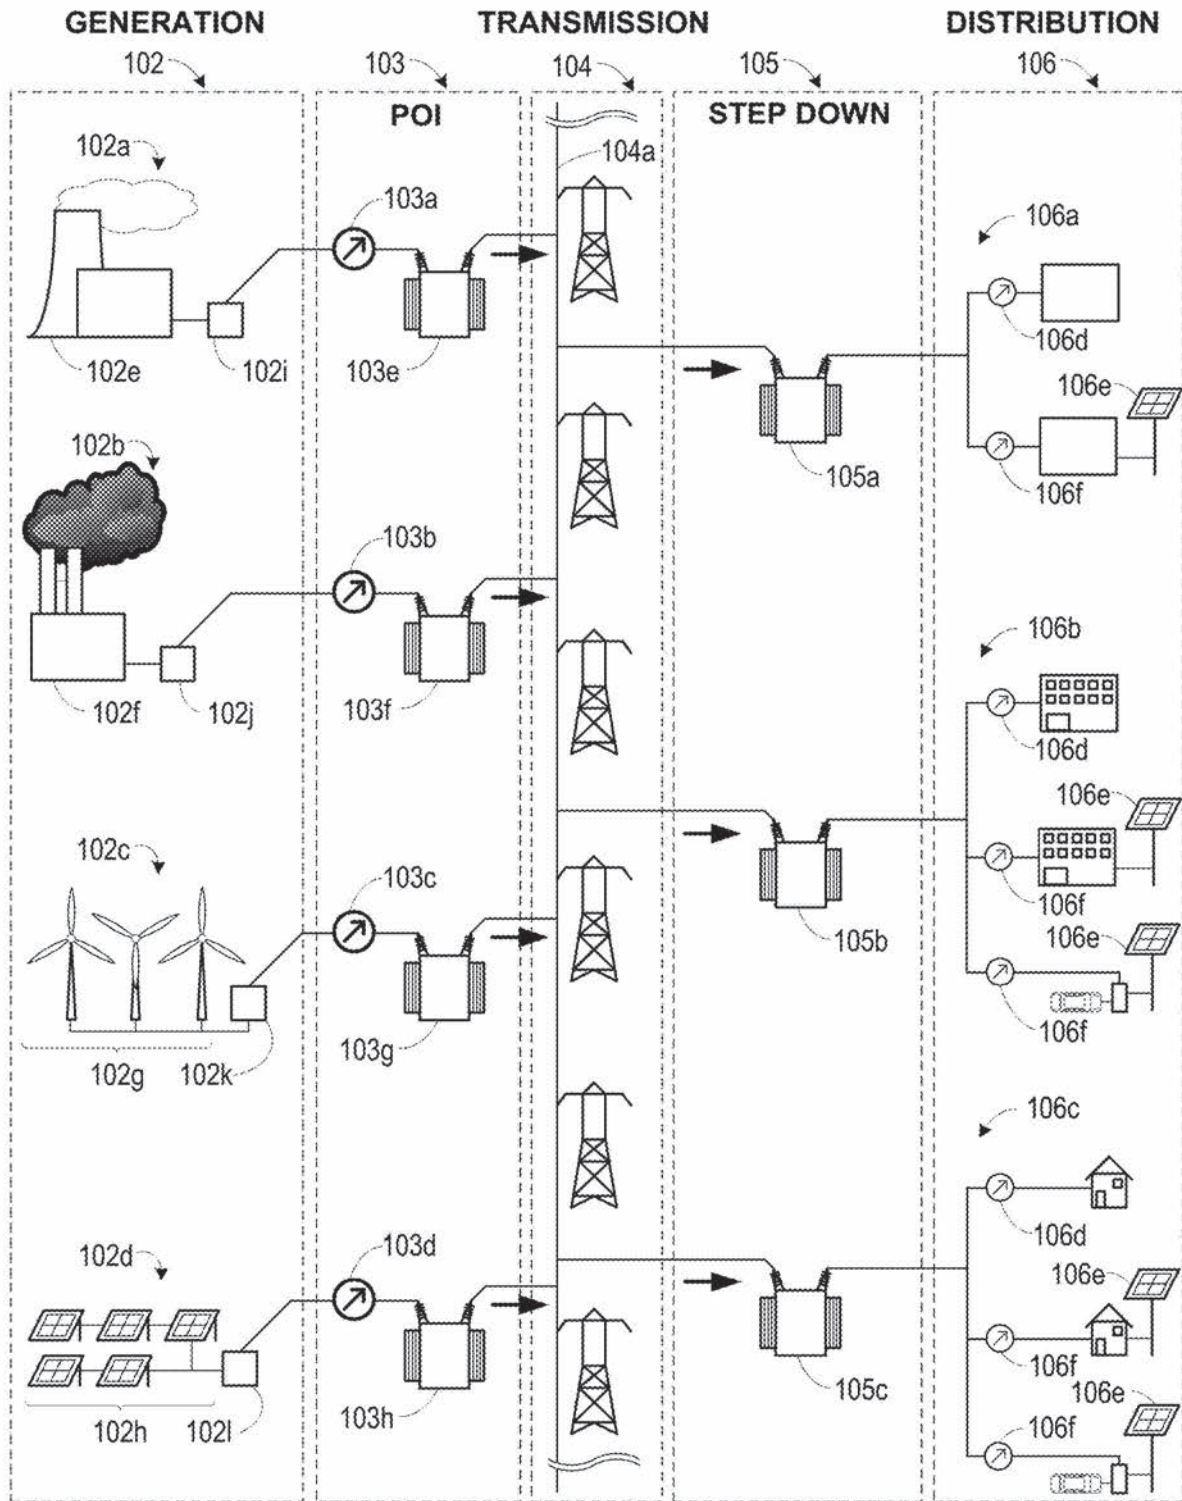
* cited by examiner

U.S. Patent

Mar. 31, 2020

Sheet 1 of 16

US 10,608,433 B1



PRIOR ART
FIGURE 1

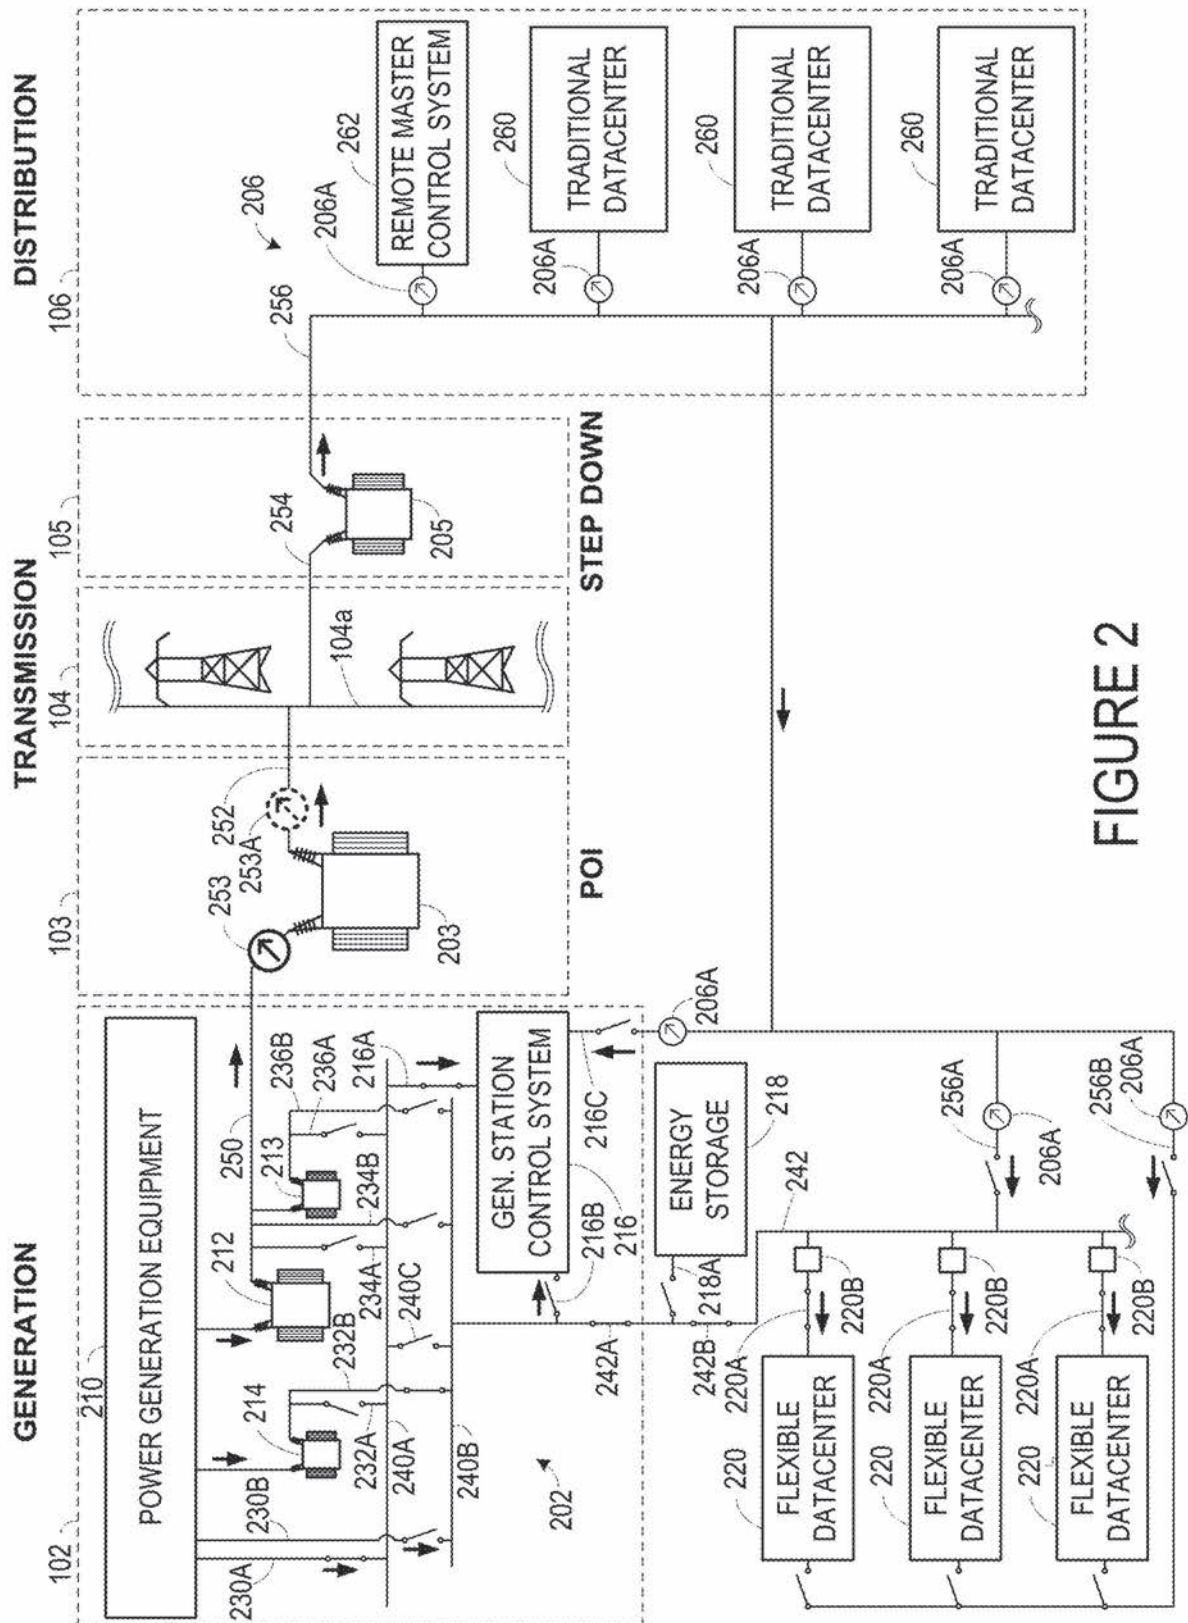


FIGURE 2

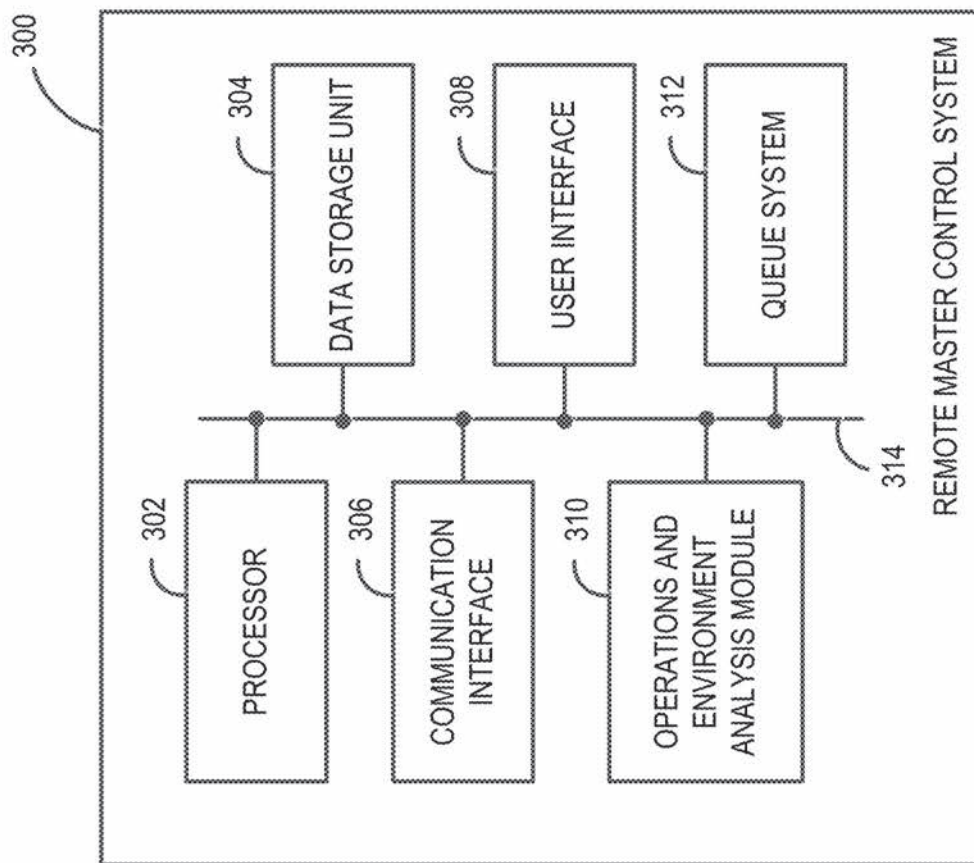


FIGURE 3

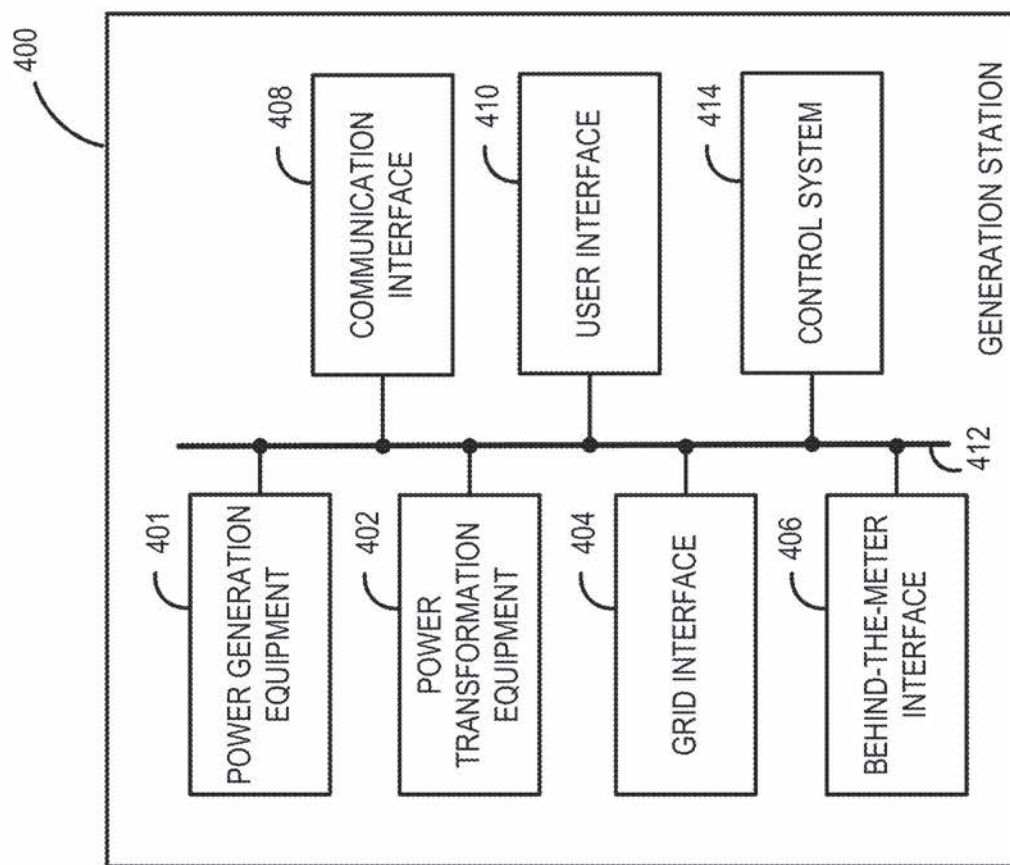


FIGURE 4

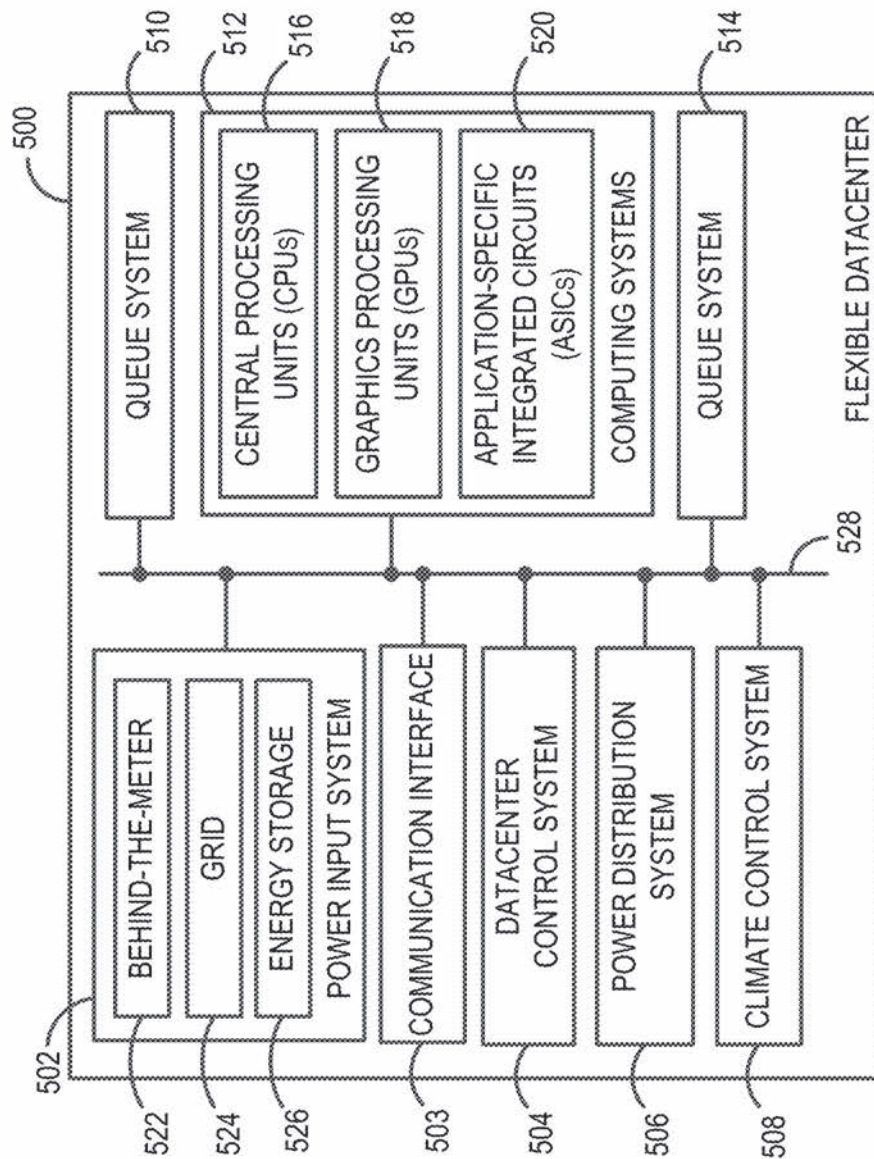


FIGURE 5

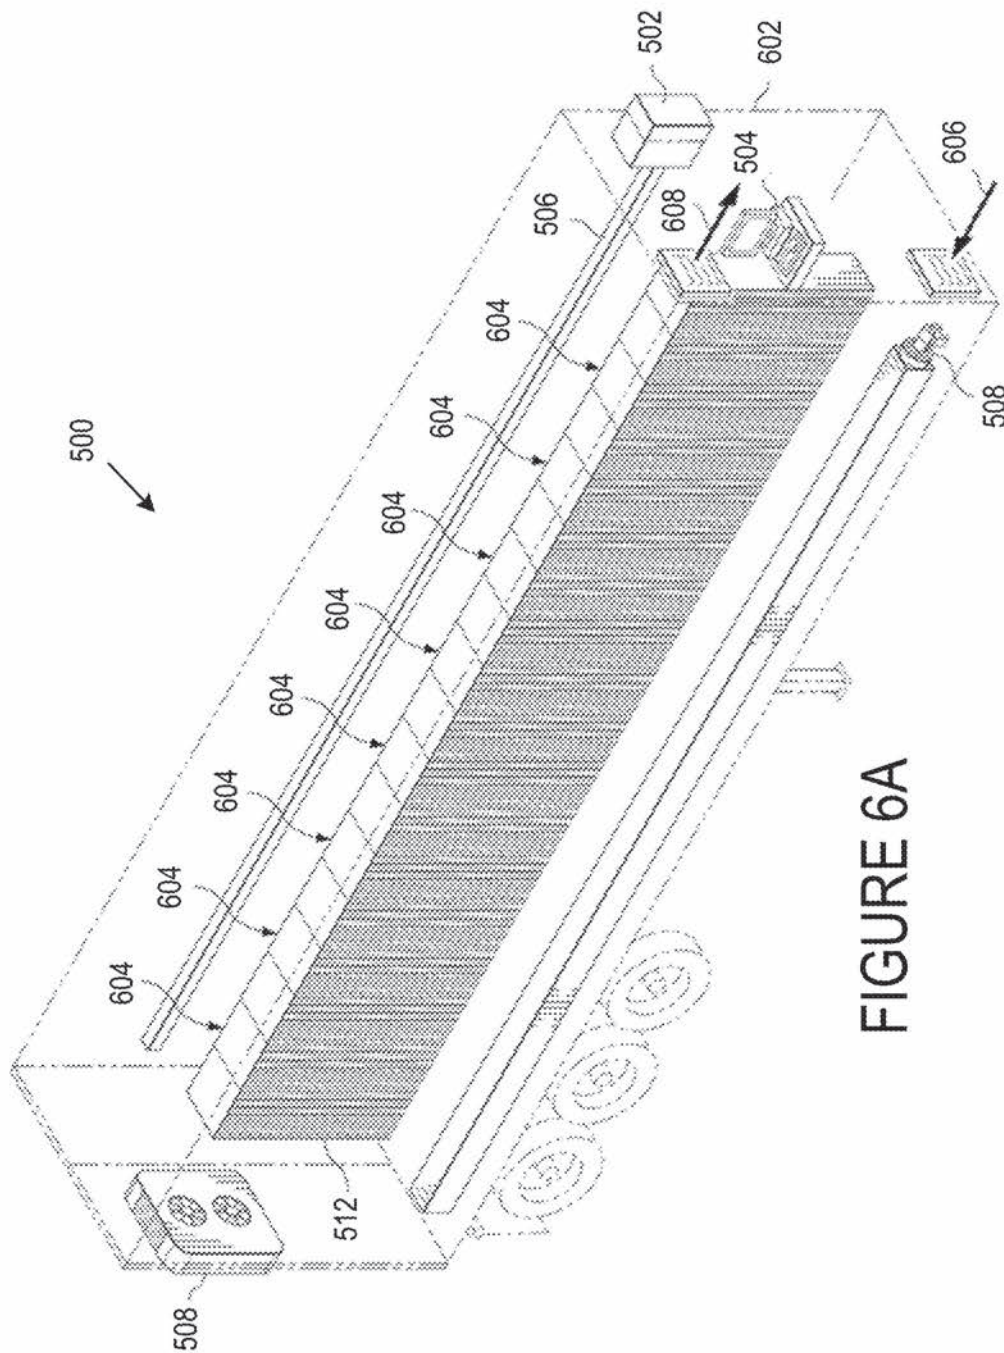


FIGURE 6A

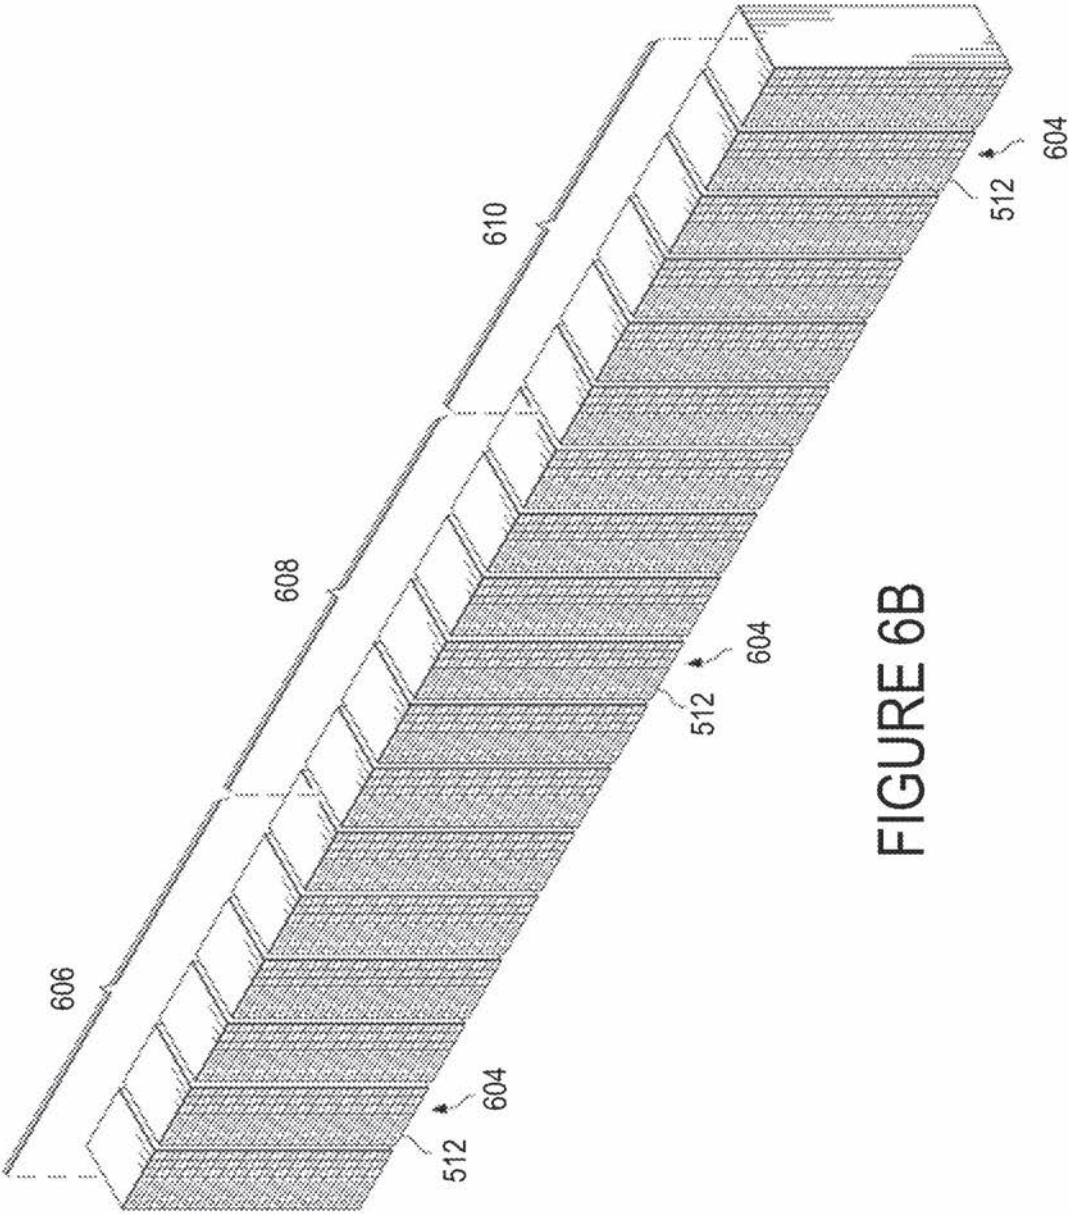


FIGURE 6B

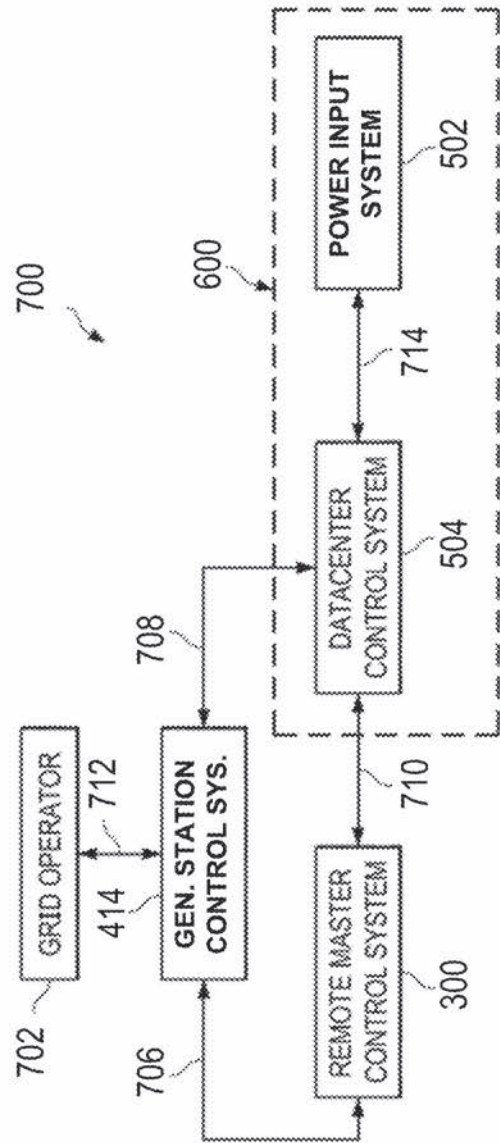
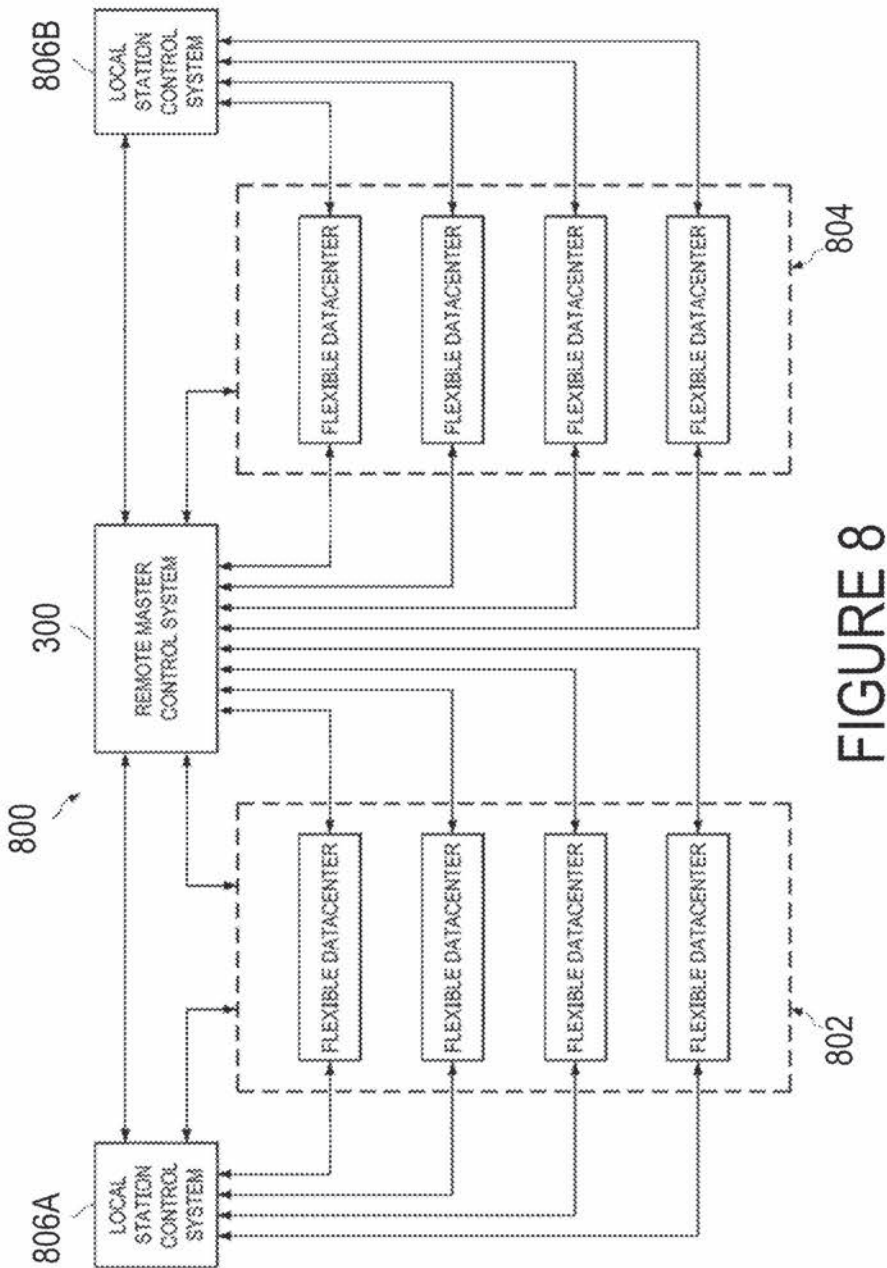


FIGURE 7



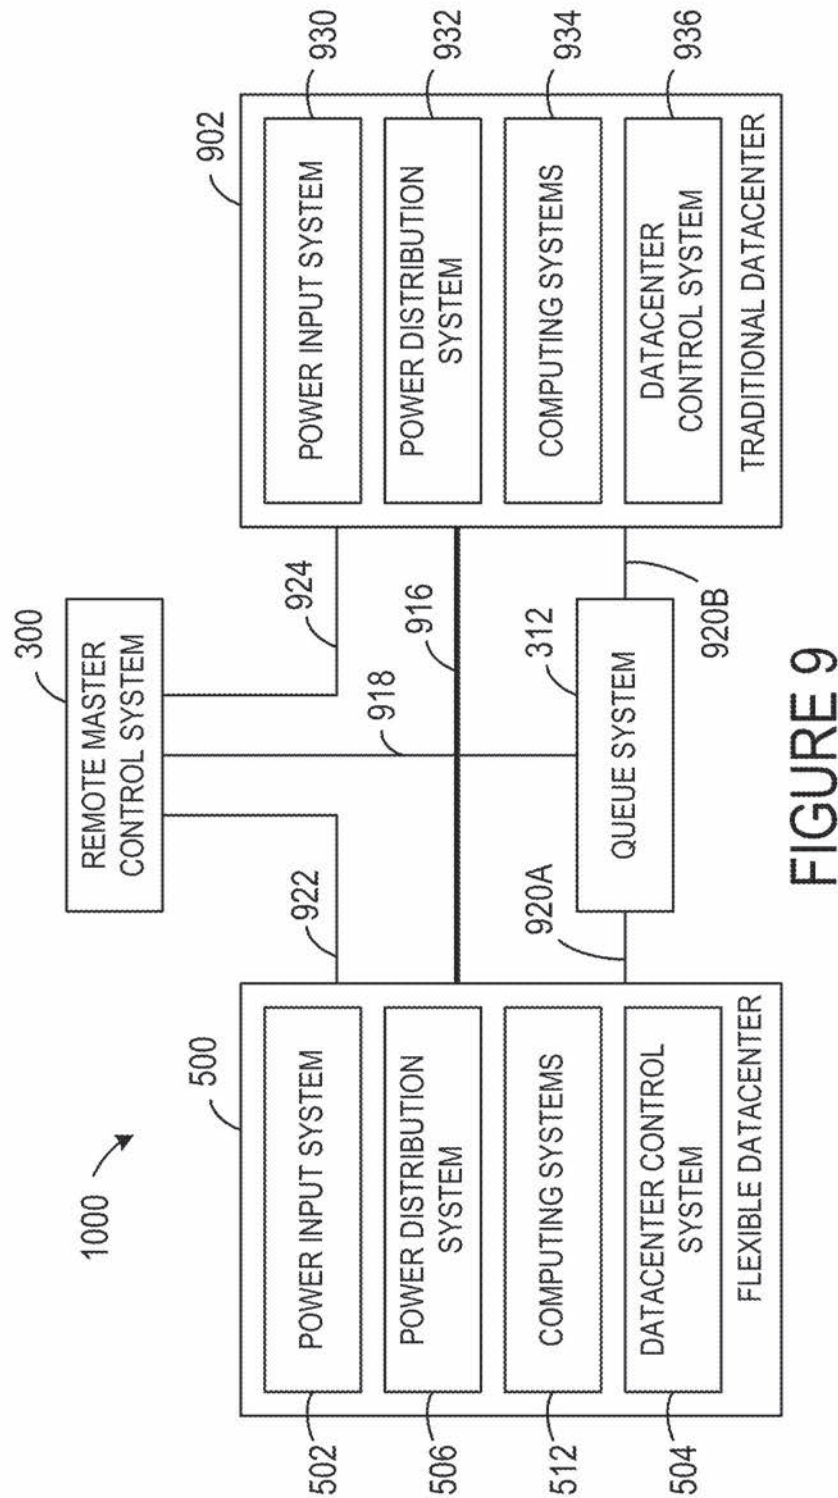


FIGURE 9

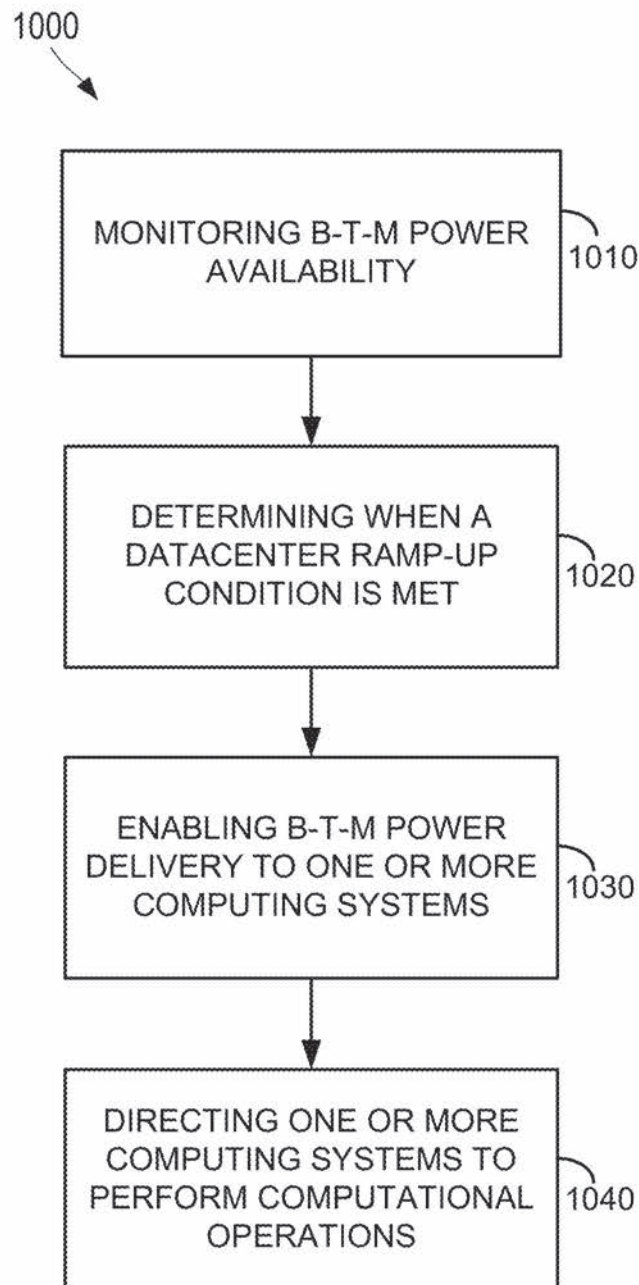


FIGURE 10A

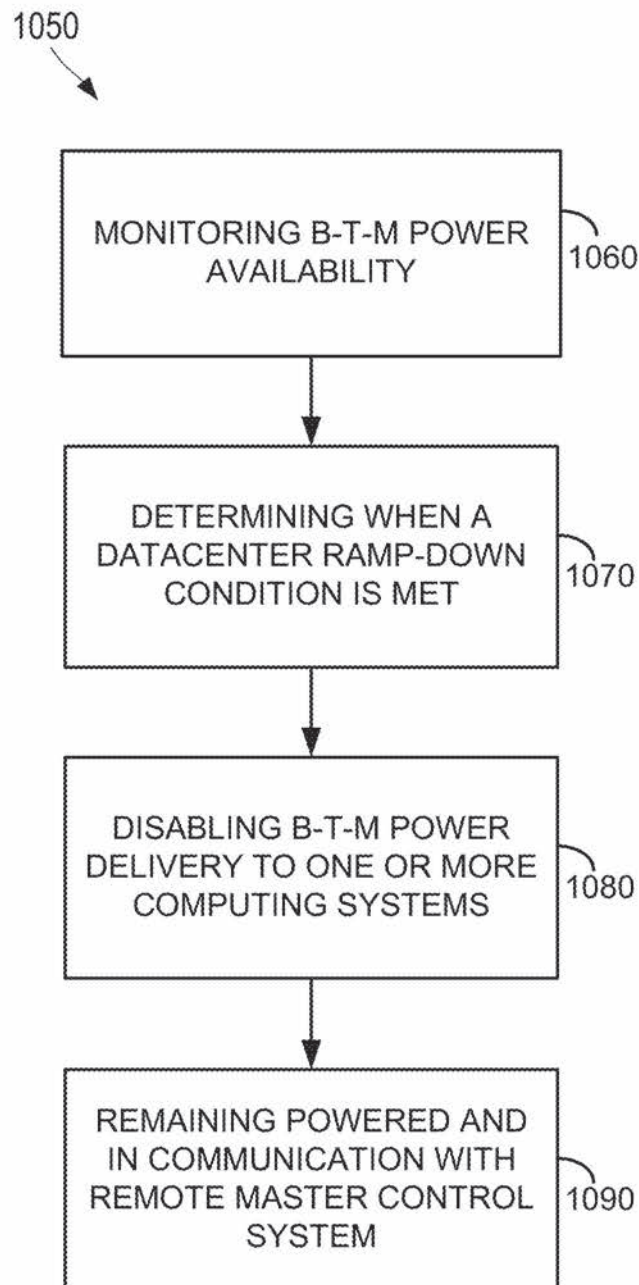


FIGURE 10B

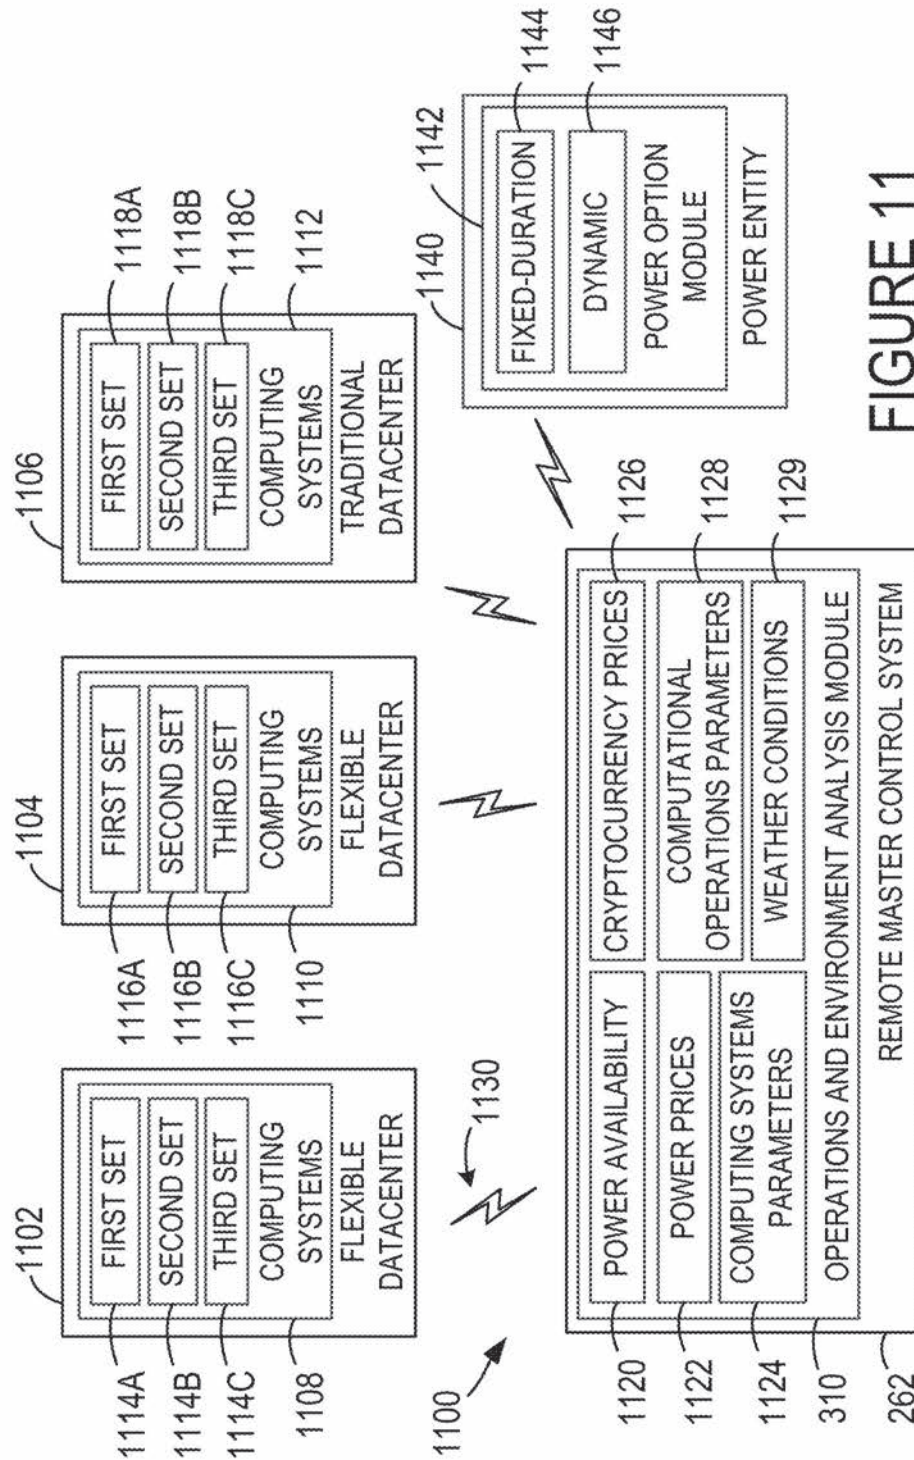


FIGURE 11

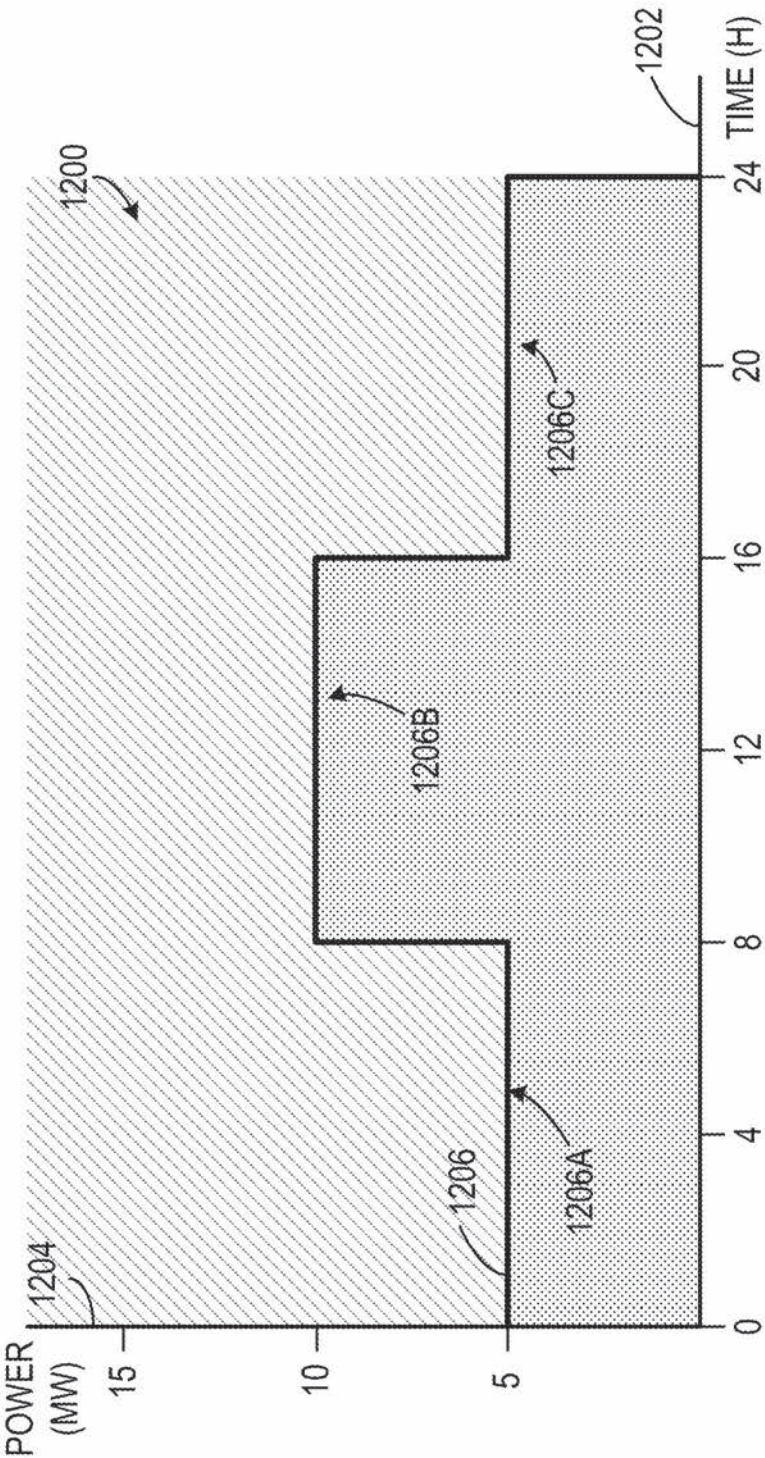


FIGURE 12

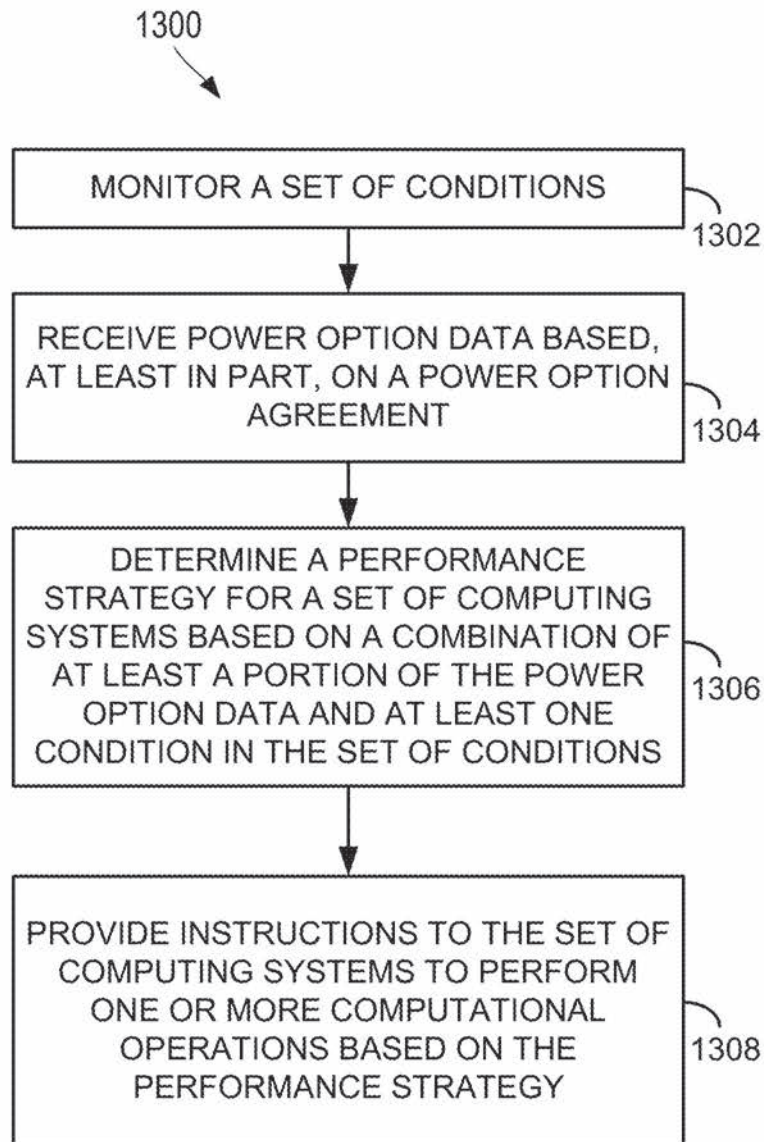


FIGURE 13

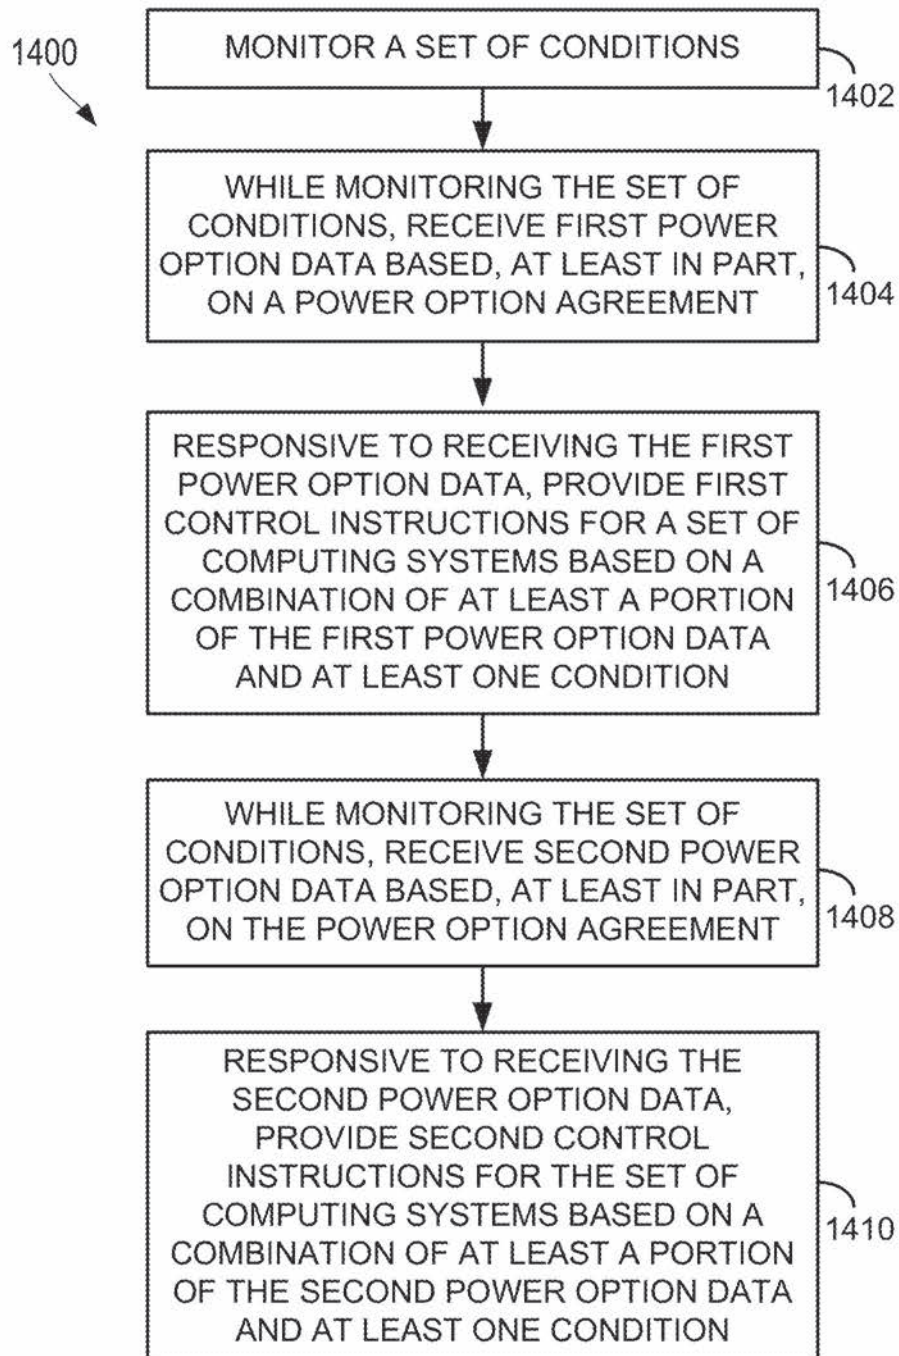


FIGURE 14

US 10,608,433 B1

1

METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Patent Application No. 62/927,119, filed Oct. 28, 2019, the entire contents of which are herein incorporated by reference.

FIELD

This specification relates to power consumption adjustments when using grid power and/or intermittent behind-the-meter power.

BACKGROUND

“Electrical grid” or “grid,” as used herein, refers to a Wide Area Synchronous Grid (also known as an Interconnection), and is a regional scale or greater electric power grid that operates at a synchronized frequency and is electrically tied together during normal system conditions. An electrical grid delivers electricity from generation stations to consumers. An electrical grid includes: (i) generation stations that produce electrical power at large scales for delivery through the grid, (ii) high voltage transmission lines that carry that power from the generation stations to demand centers, and (iii) distribution networks carry that power to individual customers.

FIG. 1 illustrates a typical electrical grid, such as a North American Interconnection or the synchronous grid of Continental Europe (formerly known as the UCTE grid). The electrical grid of FIG. 1 can be described with respect to the various segments that make up the grid.

A generation segment **102** includes one or more generation stations that produce utility-scale electricity (typically >50 MW), such as a nuclear plant **102a**, a coal plant **102b**, a wind power station (i.e., wind farm) **102c**, and/or a photovoltaic power station (i.e., a solar farm) **102d**. Generation stations are differentiated from building-mounted and other decentralized or local wind or solar power applications because they supply power at the utility level and scale (>50 MW), rather than to a local user or users. The primary purpose of generation stations is to produce power for distribution through the grid, and in exchange for payment for the supplied electricity. Each of the generation stations **102a-d** includes power generation equipment **102e-h**, respectively, typically capable of supply utility-scale power (>50 MW). For example, the power generation equipment **102g** at wind power station **102c** includes wind turbines, and the power generation equipment **102h** at photovoltaic power station **102d** includes photovoltaic panels.

Each of the generation stations **102a-d** may further include station electrical equipment **102i-1** respectively. Station electrical equipment **102i-1** are each illustrated in FIG. 1 as distinct elements for simplified illustrative purposes only and may, alternatively or additionally, be distributed throughout the power generation equipment, **102e-h**, respectively. For example, at wind power station **102c**, each wind turbine may include transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each wind turbine may be collected by distribution lines along strings of wind turbines and move through

2

collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the wind power station **102c**. Similarly, at photovoltaic power station **102d**, individual photovoltaic panels and/or arrays of photovoltaic panels may include inverters, transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each photovoltaic panel and/or array may be collected by distribution lines along the photovoltaic panels and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the photovoltaic power station **102d**.

Each generation station **102a-d** may produce AC or DC electrical current which is then typically stepped up to a higher AC voltage before leaving the respective generation station. For example, wind turbines may typically produce AC electrical energy at 600V to 700V, which may then be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102c**. As another example, photovoltaic arrays may produce DC voltage at 600V to 900V, which is then inverted to AC voltage and may be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102d**.

Upon exiting the generation segment **102**, electrical power generated at generation stations **102a-d** passes through a respective Point of Interconnection (“POI”) **103** between a generation station (e.g., **102a-d**) and the rest of the grid. A respective POI **103** represents the point of connection between a generation station’s (e.g., **102a-d**) equipment and a transmission system (e.g., transmission segment **104**) associated with electrical grid. In some cases, at the POI **103**, generated power from generation stations **102a-d** may be stepped up at transformer systems **103e-h** to high voltage scales suitable for long-distance transmission along transmission lines **104a**. Typically, the generated electrical energy leaving the POI **103** will be at 115 kV AC or above, but in some cases it may be as low as, for example, 69 kV for shorter distance transmissions along transmission lines **104a**. Each of transformer systems **103e-h** may be a single transformer or may be multiple transformers operating in parallel or series and may be co-located or located in geographically distinct locations. Each of the transformer systems **103e-h** may include substations and other links between the generation stations **102a-d** and the transmission lines **104a**.

A key aspect of the POI **103** is that this is where generation-side metering occurs. One or more utility-scale generation-side meters **103a-d** (e.g., settlement meters) are located at settlement metering points at the respective POI **103** for each generation station **102a-d**. The utility-scale generation-side meters **103a-d** measure power supplied from generation stations **102a-d** into the transmission segment **104** for eventual distribution throughout the grid.

For electricity consumption, the price consumers pay for power distributed through electric power grids is typically composed of, among other costs, Generation, Administration, and Transmission & Distribution (“T&D”) costs. T&D costs represent a significant portion of the overall price paid by consumers for electricity. These costs include capital costs (land, equipment, substations, wire, etc.), costs associated with electrical transmission losses, and operation and maintenance costs.

US 10,608,433 B1

3

For utility-scale electricity supply, operators of generation stations (e.g., 102a-d) are paid a variable market price for the amount of power the operator generates and provides to the grid, which is typically determined via a power purchase agreement (PPA) between the generation station operator and a grid operator. The amount of power the generation station operator generates and provides to the grid is measured by utility-scale generation-side meters (e.g., 103a-d) at settlement metering points. As illustrated in FIG. 1, the utility-scale generation-side meters 103a-d are shown on a low side of the transformer systems 103e-h, but they may alternatively be located within the transformer systems 103e-h or on the high side of the transformer systems 103e-h. A key aspect of a utility-scale generation-side meter is that it is able to meter the power supplied from a specific generation station into the grid. As a result, the grid operator can use that information to calculate and process payments for power supplied from the generation station to the grid. That price paid for the power supplied from the generation station is then subject to T&D costs, as well as other costs, in order to determine the price paid by consumers.

After passing through the utility-scale generation-side meters in the POI 103, the power originally generated at the generation stations 102a-d is transmitted onto and along the transmission lines 104a in the transmission segment 104. Typically, the electrical energy is transmitted as AC at 115 kV+ or above, though it may be as low as 69 kV for short transmission distances. In some cases, the transmission segment 104 may include further power conversions to aid in efficiency or stability. For example, transmission segment 104 may include high-voltage DC ("HVDC") portions (along with conversion equipment) to aid in frequency synchronization across portions of the transmission segment 104. As another example, transmission segment 104 may include transformers to step AC voltage up and then back down to aid in long distance transmission (e.g., 230 kV, 500 kV, 765 kV, etc.).

Power generated at the generation stations 104a-d is ultimately destined for use by consumers connected to the grid. Once the energy has been transmitted along the transmission segment 104, the voltage will be stepped down by transformer systems 105a-c in the step down segment 105 so that it can move into the distribution segment 106.

In the distribution segment 106, distribution networks 106a-c take power that has been stepped down from the transmission lines 104a and distribute it to local customers, such as local sub-grids (illustrated at 106a), industrial customers, including large EV charging networks (illustrated at 106b), and/or residential and retail customers, including individual EV charging stations (illustrated at 106c). Customer meters 106d, 106f measure the power used by each of the grid-connected customers in distribution networks 106a-c. Customer meters 106d are typically load meters that are unidirectional and measure power use. Some of the local customers in the distribution networks 106a-d may have local wind or solar power systems 106e owned by the customer. As discussed above, these local customer power systems 106e are decentralized and supply power directly to the customer(s). Customers with decentralized wind or solar power systems 106e may have customer meters 106f that are bidirectional or net-metering meters that can track when the local customer power systems 106e produce power in excess of the customer's use, thereby allowing the utility to provide a credit to the customer's monthly electricity bill. Customer meters 106d, 106f differ from utility-scale generation-side meters (e.g., settlement meters) in at least the following characteristics: design (electro-mechanical or electronic vs

4

current transformer), scale (typically less than 1600 amps vs. typically greater than 50 MW; typically less than 600V vs. typically greater than 14 kV), primary function (use vs. supply metering), economic purpose (credit against use vs. payment for power), and location (in a distribution network at point of use vs. at a settlement metering point at a Point of Interconnection between a generation station and a transmission line).

To maintain stability of the grid, the grid operator strives to maintain a balance between the amount of power entering the grid from generation stations (e.g., 102a-d) and the amount of grid power used by loads (e.g., customers in the distribution segment 106). In order to maintain grid stability and manage congestion, grid operators may take steps to reduce the supply of power arriving from generation stations (e.g., 102a-d) when necessary (e.g., curtailment). Particularly, grid operators may decrease the market price paid for generated power to dis-incentivize generation stations (e.g., 102a-d) from generating and supplying power to the grid. In some cases, the market price may even go negative such that generation station operators must pay for power they allow into the grid. In addition, some situations may arise where grid operators explicitly direct a generation station (e.g., 102a-d) to reduce or stop the amount of power the station is supplying to the grid.

Power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewable energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

According to the National Renewable Energy Laboratory's Technical Report TP-6A20-60983 (March 2014):

[C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission congestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

"Curtailment" today does not necessarily mean what it did in the early 2000s. Two separate changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them. Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being

US 10,608,433 B1

5

replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator’s capital costs are included in a utility’s rate base.

Curtailment may result in available energy being wasted because solar and wind operators have zero variable cost (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

SUMMARY

In an example, a system includes a set of computing systems. The set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and, while monitoring the set of conditions, receive first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The control system is further configured to provide first control instructions for the set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions responsive to receiving the first power option data. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The control system is also configured to, while monitoring the set of conditions, receive second power option data based, at least in part, on the power option

6

agreement. The second power option data specify a second minimum power threshold associated with a second time interval. Responsive to receiving the second power option data, the control system is configured to provide second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and wherein the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In another example, a method involves monitoring, at a computing system, a set of conditions, and while monitoring the set of conditions, receiving first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The method further involves, responsive to receiving the first power option data, providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The method further involves, while monitoring the set of conditions, receiving second power option data based, at least in part, on the power option agreement. The second power option data specify a second minimum power threshold associated with a second time interval. The method also involves, responsive to receiving the second power option data, providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In yet another example, a system is provided. The system includes a set of computing systems, where the set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and receive power option data based, at least in part, on a power option agreement. The power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, where each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals. The control system is further configured to, responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, where each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The control system is also configured to provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.

In a further example, non-transitory computer-readable medium is described that is configured to store instructions,

US 10,608,433 B1

7

that when executed by a computing system, causes the computing system to perform operations consistent with the method steps described above.

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 shows a typical electrical grid.

FIG. 2 shows a behind-the-meter arrangement with optional grid power, including one or more flexible datacenters, according to one or more example embodiments.

FIG. 3 shows a block diagram of a remote master control system, according to one or more example embodiments.

FIG. 4 a block diagram of a generation station, according to one or more example embodiments.

FIG. 5 shows a block diagram of a flexible datacenter, according to one or more example embodiments.

FIG. 6A shows a structural arrangement of a flexible datacenter, according to one or more example embodiments.

FIG. 6B shows a set of computing systems arranged in a straight configuration, according to one or more example embodiments.

FIG. 7 shows a control distribution system for a flexible datacenter, according to one or more example embodiments.

FIG. 8 shows a control distribution system for a fleet of flexible datacenters, according to one or more example embodiments.

FIG. 9 shows a queue distribution system for a traditional datacenter and a flexible datacenter, according to one or more example embodiments.

FIG. 10A shows a method of dynamic power consumption at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 10B shows a method of dynamic power delivery at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 11 shows a block diagram of a system for implementing power consumption adjustments based on a power option agreement, according to one or more embodiments.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments.

FIG. 13 shows a method for implementing power consumption adjustments based on a fixed-duration power option agreement, according to one or more embodiments.

FIG. 14 shows a method for implementing power consumption adjustments based on a dynamic power option agreement, according to one or more embodiments.

DETAILED DESCRIPTION

Disclosed examples will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all of the disclosed examples are shown. Different examples may be described and should not be construed as limited to the examples set forth herein.

As discussed above, the market price paid to generation stations for supplying power to the grid often fluctuates due to various factors, including the need to maintain grid stability and based on current demand and usage by connected loads in distribution networks. Due to these factors, situations can arise where generation stations are offered substantially lower prices to deter an over-supply of power to the grid. Although these situations typically exist temporarily, generation stations are sometimes forced to either sell power to the grid at the much lower prices or adjust

8

operations to decrease the amount of power generated. Furthermore, some situations may even require generation stations to incur costs in order to offload power to the grid or to shut down generation temporarily.

The volatility in the market price offered for power supplied to the grid can be especially problematic for some types of generation stations. In particular, wind farms and some other types of renewable resource power producers may lack the ability to quickly adjust operations in response to changes in the market price offered for supplying power to the grid. As a result, power generation and management at some generation stations can be inefficient, which can frequently result in power being sold to the grid at low or negative prices. In some situations, a generation station may even opt to halt power generation temporarily to avoid such unfavorable pricing. As such, the time required to halt and to restart the power generation at a generation station can reduce the generation station's ability to take advantage of rising market prices for power supplied to the grid.

Example embodiments provided herein aim to assist generation stations in managing power generation operations and avoid unfavorable power pricing situations like those described above. In particular, example embodiments may involve providing a load that is positioned behind-the-meter ("BTM") and enabling the load to utilize power received behind-the-meter at a generation station in a timely manner. As a general rule of thumb, BTM power is not subject to traditional T&D costs.

For purposes herein, a generation station is considered to be configured for the primary purpose of generating utility-scale power for supply to the electrical grid (e.g., a Wide Area Synchronous Grid or a North American Interconnect).

In one embodiment, equipment located behind-the-meter ("BTM equipment") is equipment that is electrically connected to a generation station's power generation equipment behind (i.e., prior to) the generation station's POI with an electrical grid.

In one embodiment, behind-the-meter power ("BTM power") is electrical power produced by a generation station's power generation equipment and utilized behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM equipment receives power from the generation station, but the power received by the BTM equipment from the generation station has not passed through the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM power is utilized before being metered at the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

US 10,608,433 B1

9

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that supplies power to a grid, and the BTM equipment receives power from the generation station that is not subject to T&D charges, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that supplies power to a grid, and the BTM power is not subject to T&D charges before being used by electrical equipment, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, equipment may be considered behind-the-meter if the BTM equipment receives power generated from the generation station and that received power is not routed through the electrical grid before being delivered to the BTM equipment.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station, and BTM equipment receives that generated power, and that generated power received by the BTM equipment is not routed through the electrical grid before being delivered to the BTM equipment.

For purposes herein, BTM equipment may also be referred to as a behind-the-meter load ("BTM load") when the BTM equipment is actively consuming BTM power.

Beneficially, where BTM power is not subject to traditional T&D costs, a wind farm or other type of generation station can be connected to BTM loads which can allow the generation station to selectively avoid the adverse or less-than optimal cost structure occasionally associated with supplying power to the grid by shunting generated power to the BTM load.

An arrangement that positions and connects a BTM load to a generation station can offer several advantages. In such arrangements, the generation station may selectively choose whether to supply power to the grid or to the BTM load, or both. The operator of a BTM load may pay to utilize BTM power at a cost less than that charged through a consumer meter (e.g., 106d, 1060 located at a distribution network (e.g., 106a-c) receiving power from the grid. The operator of a BTM load may additionally or alternatively charge less than the market rate to consume excess power generated at the generation station during curtailment. As a result, the generation station may direct generated power based on the "best" price that the generation station can receive during a given time frame, and/or the lowest cost the generation station may incur from negative market pricing during curtailment. The "best" price may be the highest price that the generation station may receive for its generated power during a given duration, but can also differ within embodiments and may depend on various factors, such as a prior PPA.

In one example, by having a behind-the-meter option available, a generation station may transition from supplying all generated power to the grid to supplying some or all generated power to one or more BTM loads when the market price paid for power by grid operators drops below a predefined threshold (e.g., the price that the operator of the BTM load is willing to pay the generation station for power). Thus, by having an alternative option for power consumption (i.e., one or more BTM loads), the generation station can selectively utilize the different options to maximize the price received for generated power. In addition, the generation station may also utilize a BTM load to avoid or reduce

10

the economic impact in situations when supplying power to the grid would result in the generation station incurring a net cost.

Providing BTM power to a load can also benefit the BTM load operator. A BTM load may be able to receive and utilize BTM power received from the generation station at a cost that is lower than the cost for power from the grid (e.g., at a customer meter 106d, 1060. This is primarily due to the avoidance (or significant reduction) in T&D costs and the market effects of curtailment. As indicated above, the generation station may be willing to divert generated power to the BTM load rather than supplying the grid due to changing market conditions, or during maintenance periods, or for other non-market conditions. Thus, some situations may arise where the generation station offers power to the BTM load at a price that is substantially lower than the price available on the grid. Furthermore, in some situations, the BTM load may even be able to obtain and utilize BTM power from a generation station at no cost or even at negative pricing since the generation station may rather supply the BTM load with generated power during a given time range instead of paying a higher price for the grid to take the power or modifying operations to decrease power output.

Another example of cost-effective use of BTM power is when the generation station 202 is selling power to the grid at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price the generation station 202 would have to pay to the grid power to offload generation's station 202 generated power. Advantageously, one or more flexible datacenters 220 may take the generated power behind-the-meter, thereby allowing the generation station 202 to produce and obtain the production tax credit, while selling less power to the grid at the negative price.

Another example of cost-effective behind-the-meter power is when the generation station 202 is selling power to the grid at a negative price because the grid is oversupplied and/or the generation station 202 is instructed to stand down and stop producing altogether. A grid operator may select and direct certain generation stations to go offline and stop supplying power to the grid. Advantageously, one or more flexible datacenters may be used to take power behind-the-meter, thereby allowing the generation station 202 to stop supplying power to the grid, but still stay online and make productive use of the power generated.

Another example of beneficial behind-the-meter power use is when the generation station 202 is producing power that is, with reference to the grid, unstable, out of phase, or at the wrong frequency, or the grid is already unstable, out of phase, or at the wrong frequency. A grid operator may select certain generation stations to go either offline and stop producing power, or to take corrective action with respect to the grid power stability, phase, or frequency. Advantageously, one or more flexible datacenters 220 may be used to selectively consume power behind-the-meter, thereby allowing the generation station 202 to stop providing power to the grid and/or provide corrective feedback to the grid.

Another example of beneficial behind-the-meter power use is that cost-effective behind-the-meter power availability may occur when the generation station 202 is starting up or testing. Individual equipment in the power generation equipment 210 may be routinely offline for installation, maintenance, and/or service and the individual units must be tested prior to coming online as part of overall power generation equipment 210. During such testing or maintenance time, one or more flexible datacenters may be intermittently

US 10,608,433 B1

11

powered by the one or more units of the power generation equipment **210** that are offline from the overall power generation equipment **210**.

Another example of beneficial behind-the-meter power use is that datacenter control systems at the flexible datacenters **220** may quickly ramp up and ramp down power consumption by computing systems in the flexible datacenters **220** based on power availability from the generation station **202**. For instance, if the grid requires additional power and signals the demand via a higher local price for power, the generation station **202** can supply the grid with power nearly instantly by having active flexible datacenters **220** quickly ramp down and turn off computing systems (or switch to a stored energy source), thereby reducing an active BTM load.

Another example of beneficial behind-the-meter power use is in new photovoltaic generation stations **202**. For example, it is common to design and build new photovoltaic generation stations with a surplus of power capacity to account for degradation in efficiency of the photovoltaic panels over the life of the generation stations. Excess power availability at the generation station can occur when there is excess local power generation and/or low grid demand. In high incident sunlight situations, a photovoltaic generation station **202** may generate more power than the intended capacity of generation station **202**. In such situations, a photovoltaic generation station **202** may have to take steps to protect its equipment from damage, which may include taking one or more photovoltaic panels offline or shunting their voltage to dummy loads or the ground. Advantageously, one or more flexible datacenters (e.g., the flexible datacenters **220**) may take power behind-the-meter at the Generation Station **202**, thereby allowing the generation station **202** to operate the power generation equipment **210** within operating ranges while the flexible datacenters **220** receive BTM power without transmission or distribution costs.

Thus, for at least the reasons described herein, arrangements that involves providing a BTM load as an alternative option for a generation station to direct its generated power to can serve as a mutually beneficial relationship in which both the generation station and the BTM load can economically benefit. The above-noted examples of beneficial use of BTM power are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as benefits to unutilized BTM power capacity, BTM power pricing, or BTM power consumption.

Within example embodiments described herein, various types of utility-scale power producers may operate as generation stations **202** that are capable of supplying power to one or more loads behind-the-meter. For instance, renewable energy sources (e.g., wind, solar, hydroelectric, wave, water current, tidal), fossil fuel power generation sources (coal, natural gas), and other types of power producers (e.g., nuclear power) may be positioned in an arrangement that enables the intermittent supply of generated power behind-the-meter to one or more BTM loads. One of ordinary skill in the art will recognize that the generation station **202** may vary based on an application or design in accordance with one or more example embodiments.

In addition, the particular arrangement (e.g., connections) between the generation station and one or more BTM loads can vary within examples. In one embodiment, a generation station may be positioned in an arrangement wherein the generation station selectively supplies power to the grid and/or to one or more BTM loads. As such, power cost-analysis and other factors (e.g., predicted weather condi-

12

tions, contractual obligations, etc.) may be used by the generation station, a BTM load control system, a remote master control system, or some other system or enterprise, to selectively output power to either the grid or to one or more BTM loads in a manner that maximizes revenue to the generation station. In such an arrangement, the generation station may also be able to supply both the grid and one or more BTM loads simultaneously. In some instances, the arrangement may be configured to allow dynamic manipulation of the percentage of the overall generated power that is supplied to each option at a given time. For example, in some time periods, the generation station may supply no power to the BTM load.

In addition, the type of loads that are positioned behind-the-meter can vary within example embodiments. In general, a load that is behind-the-meter may correspond to any type of load capable of receiving and utilizing power behind-the-meter from a generation station. Some examples of loads include, but are not limited to, datacenters and electric vehicle (EV) charging stations.

Preferred BTM loads are loads that can be subject to intermittent power supply because BTM power may be available intermittently. In some instances, the generation station may generate power intermittently. For example, wind power station **102c** and/or photovoltaic power station **102d** may only generate power when resource are available or favorable. Additionally or alternatively, BTM power availability at a generation station may only be available intermittently due to power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and/or operational directives from grid operators or generation station operators.

Some example embodiments of BTM loads described herein involve using one or more computing systems to serve as a BTM load at a generation station. In particular, the computing system or computing systems may receive power behind-the-meter from the generation station to perform various computational operations, such as processing or storing information, performing calculations, mining for cryptocurrencies, supporting blockchain ledgers, and/or executing applications, etc.

Multiple computing systems positioned behind-the-meter may operate as part of a "flexible" datacenter that is configured to operate only intermittently and to receive and utilize BTM power to carry out various computational operations similar to a traditional datacenter. In particular, the flexible datacenter may include computing systems and other components (e.g., support infrastructure, a control system) configured to utilize BTM power from one or more generation stations. The flexible datacenter may be configured to use particular load ramping abilities (e.g., quickly increase or decrease power usage) to effectively operate during intermittent periods of time when power is available from a generation station and supplied to the flexible datacenter behind-the-meter, such as during situations when supplying generated power to the grid is not favorable for the generation station.

In some instances, the amount of power consumed by the computing systems at a flexible datacenter can be ramped up and down quickly, and potentially with high granularity (i.e., the load can be changed in small increments if desired). This may be done based on monitored power system conditions or other information analyses as discussed herein. As recited above, this can enable a generation station to avoid negative power market pricing and to respond quickly to grid direc-

US 10,608,433 B1

13

tives. And by extension, the flexible datacenter may obtain BTM power at a price lower than the cost for power from the grid.

Various types of computing systems can provide granular power ramping. Preferably, the computing systems can perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp down or up. In some embodiments, a control system may be used to activate or de-activate one or more computing systems in an array of computing systems. For example, the control system may provide control instructions to one or more blockchain miners (e.g., a group of blockchain miners), including instructions for powering on or off, adjusting frequency of computing systems performing operations (e.g., adjusting the processing frequency), adjusting the quantity of operations being performed, and when to operate within a low power mode (if available).

Within examples, a control system may correspond to a specialized computing system or may be a computing system within a datacenter serving in the role of the control system. The location of the control system can vary within examples as well. For instance, the control system may be located at a datacenter or physically separate from the datacenter. In some examples, the control system may be part of a network of control systems that manage computational operations, power consumption, and other aspects of a fleet of datacenters. The fleet of datacenters may include one or more traditional datacenters and/or flexible datacenters.

Some embodiments may involve using one or more control systems to direct time-insensitive (e.g., interruptible) computational tasks to computational hardware, such as central processing units (CPUs) and graphics processing units (GPUs), sited behind the meter, while other hardware is sited in front of the meter (i.e., consuming metered grid power via a customer meter (e.g., 106d, 1060) and possibly remote from the behind-the-meter hardware. As such, parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, machine learning, neural network processing, queued operations, and oil and gas field simulation models, are good candidates for such interruptible computational operations.

FIG. 2 shows a behind-the-meter arrangement with optional grid-power, including one or more flexible datacenters, according to one or more example embodiments. Dark arrows illustrate a typical power delivery direction. Consistent with FIG. 1, the arrangement illustrates a generation station 202 in the generation segment 102 of a Wide-Area Synchronous Grid. The generation station 202 supplies utility-scale power (typically >50 MW) via a generation power connection 250 to the Point of Interconnection 103 between the generation station 202 and the rest of the grid. Typically, the power supplied on connection 250 may be at 34.5 kV AC, but it may be higher or lower. Depending on the voltage at connection 250 and the voltage at transmission lines 104a, a transformer system 203 may step up the power supplied from the generation station 202 to high voltage (e.g., 115 kV+AC) for transmission over connection 252 and onto transmission lines 104a of transmission segment 104. Grid power carried on the transmission segment 104 may be from generation station 202 as well as other generation stations (not shown). Also consistent with FIG. 1, grid power is consumed at one or more distribution networks, including example distribution network 206. Grid power may be taken from the transmission lines 104a via connector 254 and stepped down to distribution network

14

voltages (e.g., typically 4 kV to 26 kV AC) and sent into the distribution networks, such as distribution network 206 via distribution line 256. The power on distribution line 256 may be further stepped down (not shown) before entering individual consumer facilities such as a remote master control system 262 and/or traditional datacenters 260 via customer meters 206A, which may correspond to customer meters 106d in FIG. 1, or customer meters 106f in FIG. 1 if the respective consumer facility includes a local customer power system, such as 106e (not shown in FIG. 2).

Consistent with FIG. 1, power entering the grid from generation station 202 is metered by a utility-scale generation-side meter. A utility-scale generation-side meter 253 is shown on the low side of transformer system 203 and an alternative location is shown as 253A on the high side of transformer system 203. Both locations may be considered settlement metering points for the generation station 202 at the POI 103. Alternatively, a utility-scale generation-side meter for the generation station 202 may be located at another location consistent with the descriptions of such meters provided herein.

Generation station 202 includes power generation equipment 210, which may include, as examples, wind turbines and/or photovoltaic panels. Power generation equipment 210 may further include other electrical equipment, including but not limited to switches, busses, collectors, inverters, and power unit transformers (e.g., transformers in wind turbines).

As illustrated in FIG. 2, generation station 202 is configured to connect with BTM equipment which may function as BTM loads. In the illustrated embodiment of FIG. 2, the BTM equipment includes flexible datacenters 220. Various configurations to supply BTM power to flexible datacenters 220 within the arrangement of FIG. 2 are described herein.

In one configuration, generated power may travel from the power generation equipment 210 over one or more connectors 230A, 230B to one or more electrical busses 240A, 240B, respectively. Each of the connectors 230A, 230B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 230B is shown with an open switch, and connector 230A is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used in various embodiments when BTM power is supplied without significant power conversion to BTM loads.

In various configurations, the busses 240A and 240B may be separated by an open switch 240C or combined into a common bus by a closed switch 240C.

In another configuration, generated power may travel from the power generation equipment 210 to the high side of a local step-down transformer 214. The generated power may then travel from the low side of the local step-down transformer 214 over one or more connectors 232A, 232B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 232A, 232B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 232A is shown with an open switch, and connector 232B is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the power generation equipment 210, but the generated power must be stepped down prior to use at the BTM loads.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of a local step-up transformer 212. The generated power may

US 10,608,433 B1

15

then travel from the high side of the local step-up transformer 212 over one or more connectors 234A, 234B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 234A, 234B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 234A, 234B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of the local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 to the high side of local step-down transformer 213. The generated power may then travel from the low side of the local step-down transformer 213 over one or more connectors 236A, 236B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 236A, 236B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 236A, 236B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212, but the power must be stepped down prior to use at the BTM loads.

In one embodiment, power generated at the generation station 202 may be used to power a generation station control system 216 located at the generation station 202, when power is available. The generation station control system 216 may typically control the operation of the generation station 202. Generated power used at the generation station control system 216 may be supplied from bus 240A via connector 216A and/or from bus 240B via connector 216B. Each of the connectors 216A, 216B may be a switched connector such that power may be routed independently to 240A and/or 240B. While the generation station control system 216 can consume BTM power when powered via bus 240A or bus 240B, the BTM power taken by generation station control system 216 is insignificant in terms of rendering an economic benefit. Further, the generation station control system 216 is not configured to operate intermittently, as it generally must remain always on. Further still, the generation station control system 216 does not have the ability to quickly ramp a BTM load up or down.

In another embodiment, grid power may alternatively or additionally be used to power the generation station control system 216. As illustrated here, metered grid power from a distribution network, such as distribution network 206 for simplicity of illustration purposes only, may be used to power generation station control system 216 over connector 216C. Connector 216C may be a switched connector so that metered grid power to the generation station control system 216 can be switched on or off as needed. More commonly, metered grid power would be delivered to the generation station control system 216 via a separate distribution network (not shown), and also over a switched connector. Any such grid power delivered to the generation station control system 216 is metered by a customer meter 206A and subject to T&D costs.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202

16

through POI 103 and such grid power may power the generation station control system 216.

In some configurations, an energy storage system 218 may be connected to the generation station 202 via connector 218A, which may be a switched connector. For illustrative purposes only, connector 218A is shown with an open switch but in some embodiments it may be closed. The energy storage system 218 may be connected to bus 240A and/or bus 240B and store energy produced by the power generation equipment 210. The energy storage system may also be isolated from generation station 202 by switch 242A. In times of need, such as when the power generation equipment in an idle or off state and not generating power, the energy storage system may feed power to, for example, the flexible datacenters 220. The energy storage system may also be isolated from the flexible datacenters 220 by switch 242B.

In a preferred embodiment, as illustrated, power generation equipment 210 supplies BTM power via connector 242 to flexible datacenters 220. The BTM power used by the flexible datacenters 220 was generated by the generation station 202 and did not pass through the POI 103 or utility-scale generation-side meter 253, and is not subject to T&D charges. Power received at the flexible datacenters 220 may be received through respective power input connectors 220A. Each of the respective connectors 220A may be a switched connector that can electrically isolate the respective flexible datacenter 220 from the connector 242. Power equipment 220B may be arranged between the flexible datacenters 220 and the connector 242. The power equipment 220B may include, but is not limited to, power conditioners, unit transformers, inverters, and isolation equipment. As illustrated, each flexible datacenter 220 may be served by a respective power equipment 220B. However, in another embodiment, one power equipment 220B may serve multiple flexible datacenter 220.

In one embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter and electrically connected to the power generation equipment 210 behind (i.e., prior to) the generation station's POI 103 with the rest of the electrical grid.

In one embodiment, BTM power produced by the power generation equipment 210 is utilized by the flexible datacenters 220 behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as the flexible datacenters 220 are electrically connected to the generation station 202, and generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), and the flexible datacenters 220 receive power from the generation station 202, but the power received by the flexible datacenters 220 from the generation station 202 has not passed through a utility-scale generation-side meter. In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where power produced at the generation station 202 is subject to metering by utility-scale generation-side meter

US 10,608,433 B1

17

253 (or 253A, or another utility-scale generation-side meter), but the BTM power supplied to the flexible datacenters 220 is utilized before being metered at the utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter). In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as they are electrically connected to the generation station 202 that supplies power to the grid, and the flexible datacenters 220 receive power from the generation station 202 that is not subject to T&D charges, but power otherwise received from the grid that is supplied by the generation station 202 is subject to T&D charges.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 that supplies power to a grid, and the generated power is not subject to T&D charges before being used by flexible datacenters 220, but power otherwise received from the connected grid is subject to T&D charges.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter because they receive power generated from the generation station 202 intended for the grid, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 for distribution to the grid, and the flexible datacenters 220 receive that power, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, metered grid power may alternatively or additionally be used to power one or more of the flexible datacenters 220, or a portion within one or more of the flexible datacenters 220. As illustrated here for simplicity, metered grid power from a distribution network, such as distribution network 206, may be used to power one or more flexible datacenters 220 over connector 256A and/or 256B. Each of connector 256A and/or 256B may be a switched connector so that metered grid power to the flexible datacenters 220 can be switched on or off as needed. More commonly, metered grid power would be delivered to the flexible datacenters 220 via a separate distribution network (not shown), and also over switched connectors. Any such grid power delivered to the flexible datacenters 220 is metered by customer meters 206A and subject to T&D costs. In one embodiment, connector 256B may supply metered grid power to a portion of one or more flexible datacenters 220. For example, connector 256B may supply metered grid power to control and/or communication systems for the flexible datacenters 220 that need constant power and cannot be subject to intermittent BTM power. Connector 242 may supply solely BTM power from the generation station 202 to high power demand computing systems within the flexible datacenters 220, in which case at least a portion of each flexible datacenters 220 so connected is operating as a BTM load. In another embodiment, connector 256A and/or 256B may supply all power used at one or more of the flexible

18

datacenters 220, in which case each of the flexible datacenters 220 so connected would not be operating as a BTM load.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the flexible datacenters 220.

The flexible datacenters 220 are shown in an example arrangement relative to the generation station 202. Particularly, generated power from the generation station 202 may be supplied to the flexible datacenters 220 through a series of connectors and/or busses (e.g., 232B, 240B, 242, 220A). As illustrated, in other embodiments, connectors between the power generation equipment 210 and other components may be switched open or closed, allowing other pathways for power transfer between the power generation equipment 210 and components, including the flexible datacenters 220. Additionally, the connector arrangement shown is illustrative only and other circuit arrangements are contemplated within the scope of supplying BTM power to a BTM load at generation station 202. For example, there may be more or fewer transformers, or one or more of transformers 212, 213, 214 may be transformer systems with multiple steppings and/or may include additional power equipment including but not limited to power conditioners, filters, switches, inverters, and/or AC/DC-DC/AC isolators. As another example, metered grid power connections to flexible datacenters 220 are shown via both 256A and 256B; however, a single connection may connect one or more flexible datacenters 220 (or power equipment 220B) to metered grid power and the one or more flexible datacenters 220 (or power equipment 220B) may include switching apparatus to direct BTM power and/or metered grid power to control systems, communication systems, and/or computing systems as desired.

In some examples, BTM power may arrive at the flexible datacenters 220 in a three-phase AC format. As such, power equipment (e.g., power equipment 220B) at one or more of the flexible datacenters 220 may enable each flexible datacenter 220 to use one or more phases of the power. For instance, the flexible datacenters 220 may utilize power equipment (e.g., power equipment 220B, or alternatively or additionally power equipment that is part of the flexible datacenter 220) to convert BTM power received from the generation station 202 for use at computing systems at each flexible datacenter 220. In other examples, the BTM power may arrive at one or more of the flexible datacenters 220 as DC power. As such, the flexible datacenters 220 may use the DC power to power computing systems. In some such examples, the DC power may be routed through a DC-to-DC converter that is part of power equipment 220B and/or flexible datacenter 220.

In some configurations, a flexible datacenter 220 may be arranged to only have access to power received behind-the-meter from a generation station 202. In the arrangement of FIG. 2, the flexible datacenters 220 may be arranged only with a connection to the generation station 202 and depend solely on power received behind-the-meter from the generation station 202. Alternatively or additionally, the flexible datacenters 220 may receive power from energy storage system 218.

In some configurations, one or more of the flexible datacenters 220 can be arranged to have connections to multiple sources that are capable of supplying power to a flexible datacenter 220. To illustrate a first example, the flexible datacenters 220 are shown connected to connector 242, which can be connected or disconnected via switches to

US 10,608,433 B1

19

the energy storage system 218 via connector 218A, the generation station 202 via bus 240B, and grid power via metered connector 256A. In one embodiment, the flexible datacenters 220 may selectively use power received behind-the-meter from the generation station 202, stored power supplied by the energy storage system 218, and/or grid power. For instance, flexible datacenters 220 may use power stored in the energy storage system 218 when costs for using power supplied behind-the-meter from the generation station 202 are disadvantageous. By having access to the energy storage system 218 available, the flexible datacenters 220 may use the stored power and allow the generation station 202 to subsequently refill the energy storage system 218 when cost for power behind-the-meter is low. Alternatively, the flexible datacenters 220 may use power from multiple sources simultaneously to power different components (e.g., a first set and a second set of computing systems). Thus, the flexible datacenters 220 may leverage the multiple connections in a manner that can reduce the cost for power used by the computing systems at the flexible datacenters 220. The flexible datacenters 220 control system or the remote master control system 262 may monitor power conditions and other factors to determine whether the flexible datacenters 220 should use power from either the generation station 202, grid power, the energy storage system 218, none of the sources, or a subset of sources during a given time range. Other arrangements are possible as well. For example, the arrangement of FIG. 2 illustrates each flexible datacenter 220 as connected via a single connector 242 to energy storage system 218, generation station 202, and metered grid power via 256A. However, one or more flexible datacenters 220 may have independent switched connections to each energy source, allowing the one or more flexible datacenters 220 to operate from different energy sources than other flexible datacenters 220 at the same time.

The selection of which power source to use at a flexible datacenter (e.g., the flexible datacenters 220) or another type of BTM load can change based on various factors, such as the cost and availability of power from both sources, the type of computing systems using the power at the flexible datacenters 220 (e.g., some systems may require a reliable source of power for a long period), the nature of the computational operations being performed at the flexible datacenters 220 (e.g., a high priority task may require immediate completion regardless of cost), and temperature and weather conditions, among other possible factors. As such, a datacenter control system at the flexible datacenters 220, the remote master control system 262, or another entity (e.g., an operator at the generation station 202) may also influence and/or determine the source of power that the flexible datacenters 220 use at a given time to complete computational operations.

In some example embodiments, the flexible datacenters 220 may use power from the different sources to serve different purposes. For example, the flexible datacenters 220 may use metered power from grid power to power one or more systems at the flexible datacenters 220 that are configured to be always-on (or almost always on), such as a control and/or communication system and/or one or more computing systems (e.g., a set of computing systems performing highly important computational operations). The flexible datacenters 220 may use BTM power to power other components within the flexible datacenters 220, such as one or more computing systems that perform less critical computational operations.

In some examples, one or more flexible datacenters 220 may be deployed at the generation station 202. In other

20

examples, flexible datacenters 220 may be deployed at a location geographically remote from the generation station 202, while still maintaining a BTM power connection to the generation station 202.

In another example arrangement, the generation station 202 may be connected to a first BTM load (e.g., a flexible datacenter 220) and may supply power to additional BTM loads via connections between the first BTM load and the additional BTM loads (e.g., a connection between a flexible datacenter 220 and another flexible datacenter 220).

The arrangement in FIG. 2, and components included therein, are for non-limiting illustration purposes and other arrangements are contemplated in examples. For instance, in another example embodiment, the arrangement of FIG. 2 may include more or fewer components, such as more BTM loads, different connections between power sources and loads, and/or a different number of datacenters. In addition, some examples may involve one or more components within the arrangement of FIG. 2 being combined or further divided.

Within the arrangement of FIG. 2, a control system, such as the remote master control system 262 or another component (e.g., a control system associated with the grid operator, the generation station control system 216, or a datacenter control system associated with a traditional datacenter or one or more flexible datacenters) may use information to efficiently manage various operations of some of the components within the arrangement of FIG. 2. For example, the remote master control system 262 or another component may manage distribution and execution of computational operations at one or more traditional datacenters 260 and/or flexible datacenters 220 via one or more information-processing algorithms. These algorithms may utilize past and current information in real-time to manage operations of the different components. These algorithms may also make some predictions based on past trends and information analysis. In some examples, multiple computing systems may operate as a network to process information.

Information used to make decisions may include economic and/or power-related information, such as monitored power system conditions. Monitored power system conditions may include one or more of excess power generation at a generation station 202, excess power at a generation station 202 that a connected grid cannot receive, power generation at a generation station 202 subject to economic curtailment, power generation at a generation station 202 subject to reliability curtailment, power generation at a generation station 202 subject to power factor correction, low power generation at a generation station 202, start up conditions at a generation station 202, transient power generation conditions at a generation station 202, or testing conditions where there is an economic advantage to using behind-the-meter power generation at a generation station 202. These different monitored power system conditions can be weighted differently during processing and analysis.

In some examples, the information can include the cost for power from available sources (e.g., BTM power at the generation station 202 versus metered grid power) to enable comparisons to be made which power source costs less. In some instances, the information may include historic prices for power to enable the remote master control system 262 or another system to predict potential future prices in similar situations (e.g., the cost of power tends to trend upwards for grid power during warmer weather and peak-use hours). The information may also indicate the availability of power from the various sources (e.g., BTM power at the generation

US 10,608,433 B1

21

station 262, the energy storage system 218 at the generation station 262, and/or metered grid power).

In addition, the information may also include other data, including information associated with operations at components within the arrangement. For instance, the information may include data associated with performance of operations at the flexible datacenters 220 and the traditional datacenters 260, such as the number of computational tasks currently being performed, the types of tasks being performed (e.g., type of computational operation, time-sensitivity, etc.), the number, types, and capabilities of available computing systems, the amount of computational tasks awaiting performance, and the types of computing systems at one or more datacenters, among others. The information may also include data specifying the conditions at one or more datacenters (e.g., whether or not the temperatures are in a desired range, the amount of power available within an energy storage system such as 218), the amount of computational tasks awaiting performance in the queue of one or more of the datacenters, and the identities of the entities associated with the computational operations at one or more of the datacenters. Entities associated with computational operations may be, for example, owners of the datacenters, customers who purchase computational time at the datacenters, or other entities.

The information used by the remote master control system 262 or another component may include data associated with the computational operations to be performed, such as deadlines, priorities (e.g., high vs. low priority tasks), cost to perform based on required computing systems, the optimal computing systems (e.g., CPU vs GPU vs ASIC; processing unit capabilities, speeds, or frequencies, or instructional sets executable by the processing units) for performing each requested computational task, and prices each entity (e.g., company) is willing to pay for computational operations to be performed or otherwise supported via computing systems at a traditional datacenter 260 or a flexible datacenter 220, among others. In addition, the information may also include other data (e.g., weather conditions at locations of datacenters or power sources, any emergencies associated with a datacenter or power source, or the current value of bids associated with an auction for computational tasks).

The information may be updated in-real time and used to make the different operational decisions within the arrangement of FIG. 2. For instance, the information may help a component (e.g., the remote master control system 262 or a control system at a flexible datacenter 220) determine when to ramp up or ramp down power use at a flexible datacenter 220 or when to switch one or more computing systems at a flexible datacenter 220 into a low power mode or to operate at a different frequency, among other operational adjustments. The information can additionally or alternatively help a component within the arrangement of FIG. 2 to determine when to transfer computational operations between computing systems or between datacenters based on various factors. In some instances, the information may also be used to determine when to temporarily stop performing a computational operation or when to perform a computational operation at multiple sites for redundancy or other reasons. The information may further be used to determine when to accept new computational operations from entities or when to temporarily suspend accepting new tasks to be performed due to lack of computing system availability.

The remote master control system 262 represents a computing system that is capable of obtaining, managing, and using the information described above to manage and over-

22

see one or more operations within the arrangement of FIG. 2. As such, the remote master control system 262 may be one or more computing systems configured to process all, or a subset of, the information described above, such as power, environment, computational characterization, and economic factors to assist with the distribution and execution of computing operations among one or more datacenters. For instance, the remote master control system 262 may be configured to obtain and delegate computational operations among one or more datacenters based on a weighted analysis of a variety of factors, including one or more of the cost and availability of power, the types and availability of the computing systems at each datacenter, current and predicted weather conditions at the different locations of flexible datacenters (e.g., flexible datacenters 220) and generation stations (e.g., generation stations 202), levels of power storage available at one or more energy storage systems (e.g., energy storage system 218), and deadlines and other attributes associated with particular computational operations, among other possible factors. As such, the analysis of information performed by the remote master control system 262 may vary within examples. For instance, the remote master control system 262 may use real-time information to determine whether or not to route a computational operation to a particular flexible datacenter (e.g., a flexible datacenter 220) or to transition a computational operation between datacenters (e.g., from traditional datacenter 260 to a flexible datacenter 220).

As shown in FIG. 2, the generation station 202 may be able to supply power to the grid and/or BTM loads such as flexible datacenters 220. With such a configuration, the generation station 202 may selectively provide power to the BTM loads and/or the grid based on economic and power availability considerations. For example, the generation station 202 may supply power to the grid when the price paid for the power exceeds a particular threshold (e.g., the power price offered by operators of the flexible datacenters 220). In some instances, the operator of a flexible datacenter and the operator of a generation station capable of supplying BTM power to the flexible datacenter may utilize a predefined arrangement (e.g., a contract) that specifies a duration and/or price range when the generation station may supply power to the flexible datacenter.

The remote master control system 262 may be capable of directing one or more flexible datacenters 220 to ramp-up or ramp-down to desired power consumption levels, and/or to control cooperative action of multiple flexible datacenters by determining how to power each individual flexible datacenter 220 in accordance with operational directives.

The configuration of the remote master control system 262 can vary within examples as further discussed with respect to FIGS. 2, 3, and 7-9. The remote master control system 262 may operate as a single computing system or may involve a network of computing systems. Preferably, the remote master control system 262 is implemented across one or more servers in a fault-tolerant operating environment that ensures continuous uptime and connectivity by virtue of its distributed nature. Alternatively, although the remote master control system 262 is shown as a physically separate component arrangement for FIG. 2, the remote master control system 262 may be combined with another component in other embodiments. To illustrate an example, the remote master control system 262 may operate as part of a flexible datacenter (e.g., a computing system or a datacenter control system of the flexible datacenter 220), includ-

US 10,608,433 B1

23

ing sharing components with a flexible datacenter, sharing power with a flexible datacenter, and/or being co-located with a flexible datacenter.

In addition, the remote master control system 262 may communicate with components within the arrangement of FIG. 2 using various communication technologies, including wired and wireless communication technologies. For instance, the remote master control system 262 may use wired (not illustrated) or wireless communication to communicate with datacenter control systems or other computing systems at the flexible datacenters 220 and the traditional datacenters 260. The remote master control system 262 may also communicate with entities inside or outside the arrangement of FIG. 2 and other components within the arrangement of FIG. 2 via wired or wireless communication. For instance, the remote master control system 262 may use wireless communication to obtain computational operations from entities seeking support for the computational operations at one or more datacenters in exchange for payment. The remote master control system 262 may communicate directly with the entities or may obtain the computational operations from the traditional datacenters 260. For instance, an entity may submit jobs (e.g., computational operations) to one or more traditional datacenters 260. The remote master control system 262 may determine that transferring one or more of the computational operations to a flexible datacenter 220 may better support the transferred computational operations. For example, the remote master control system 262 may determine that the transfer may enable the computational operations to be completed quicker and/or at a lower cost. In some examples, the remote master control system 262 may communicate with the entity to obtain approval prior to transferring the one or more computational operations.

The remote master control system 262 may also communicate with grid operators and/or an operator of generation station 202 to help determine power management strategies when distributing computational operations across the various datacenters. In addition, the remote master control system 262 may communicate with other sources, such as weather prediction systems, historical and current power price databases, and auction systems, etc.

In further examples, the remote master control system 262 or another computing system within the arrangement of FIG. 2 may use wired or wireless communication to submit bids within an auction that involves a bidder (e.g., the highest bid) obtaining computational operations or other tasks to be performed. Particularly, the remote master control system 262 may use the information discussed above to develop bids to obtain computing operations for performance at available computing systems at flexible datacenters (e.g., flexible datacenters 220).

In the example arrangement shown in FIG. 2, the flexible datacenters 220 represent example loads that can receive power behind-the-meter from the generation station 202. In such a configuration, the flexible datacenters 220 may obtain and utilize power behind-the-meter from the generation station 202 to perform various computational operations. Performance of a computational operation may involve one or more computing systems providing resources useful in the computational operation. For instance, the flexible datacenters 220 may include one or more computing systems configured to store information, perform calculations and/or parallel processes, perform simulations, mine cryptocurrencies, and execute applications, among other potential tasks. The computing systems can be specialized or generic and can be arranged at each flexible datacenter 220 in a variety

24

of ways (e.g., straight configuration, zig-zag configuration) as further discussed with respect to FIGS. 6A, 6B. Furthermore, although the example arrangement illustrated in FIG. 2 shows configurations where flexible datacenters 220 serve as BTM loads, other types of loads can be used as BTM loads within examples.

The arrangement of FIG. 2 includes the traditional datacenters 260 coupled to metered grid power. The traditional datacenters 260 using metered grid power to provide computational resources to support computational operations. One or more enterprises may assign computational operations to the traditional datacenters 260 with expectations that the datacenters reliably provide resources without interruption (i.e., non-intermittently) to support the computational operations, such as processing abilities, networking, and/or volatile storage. Similarly, one or more enterprises may also request computational operations to be performed by the flexible datacenters 220. The flexible datacenters 220 differ from the traditional datacenters 260 in that the flexible datacenters 220 are arranged and/or configured to be connected to BTM power, are expected to operate intermittently, and are expected to ramp load (and thus computational capability) up or down regularly in response to control directives. In some examples, the flexible datacenters 220 and the traditional datacenters 260 may have similar configurations and may only differ based on the source(s) of power relied upon to power internal computing systems. Preferably, however, the flexible datacenters 220 include particular fast load ramping abilities (e.g., quickly increase or decrease power usage) and are intended and designed to effectively operate during intermittent periods of time.

FIG. 3 shows a block diagram of the remote master control system 300 according to one or more example embodiments. Remote master control system 262 may take the form of remote master control system 300, or may include less than all components in remote master control system 300, different components than in remote master control system 300, and/or more components than in remote master control system 300.

The remote master control system 300 may perform one or more operations described herein and may include a processor 302, a data storage unit 304, a communication interface 306, a user interface 308, an operations and environment analysis module 310, and a queue system 312. In other examples, the remote master control system 300 may include more or fewer components in other possible arrangements.

As shown in FIG. 3, the various components of the remote master control system 300 can be connected via one or more connection mechanisms (e.g., a connection mechanism 314). In this disclosure, the term "connection mechanism" means a mechanism that facilitates communication between two or more devices, systems, components, or other entities. For instance, a connection mechanism can be a simple mechanism, such as a cable, PCB trace, or system bus, or a relatively complex mechanism, such as a packet-based communication network (e.g., LAN, WAN, and/or the Internet). In some instances, a connection mechanism can include a non-tangible medium (e.g., where the connection is wireless).

As part of the arrangement of FIG. 2, the remote master control system 300 (corresponding to remote master control system 262) may perform a variety of operations, such as management and distribution of computational operations among datacenters, monitoring operational, economic, and environment conditions, and power management. For instance, the remote master control system 300 may obtain

US 10,608,433 B1

25

computational operations from one or more enterprises for performance at one or more datacenters. The remote master control system 300 may subsequently use information to distribute and assign the computational operations to one or more datacenters (e.g., the flexible datacenters 220) that have the resources (e.g., particular types of computing systems and available power) available to complete the computational operations. In some examples, the remote master control system 300 may assign all incoming computational operation requests to the queue system 312 and subsequently assign the queued requests to computing systems based on an analysis of current market and power conditions.

Although the remote master control system 300 is shown as a single entity, a network of computing systems may perform the operations of the remote master control system 300 in some examples. For example, the remote master control system 300 may exist in the form of computing systems (e.g., datacenter control systems) distributed across multiple datacenters.

The remote master control system 300 may include one or more processors 302. As such, the processor 302 may represent one or more general-purpose processors (e.g., a microprocessor) and/or one or more special-purpose processors (e.g., a digital signal processor (DSP)). In some examples, the processor 302 may include a combination of processors within examples. The processor 302 may perform operations, including processing data received from the other components within the arrangement of FIG. 2 and data obtained from external sources, including information such as weather forecasting systems, power market price systems, and other types of sources or databases.

The data storage unit 304 may include one or more volatile, non-volatile, removable, and/or non-removable storage components, such as magnetic, optical, or flash storage, and/or can be integrated in whole or in part with the processor 302. As such, the data storage unit 304 may take the form of a non-transitory computer-readable storage medium, having stored thereon program instructions (e.g., compiled or non-compiled program logic and/or machine code) that, when executed by the processor 302, cause the remote master control system 300 to perform one or more acts and/or functions, such as those described in this disclosure. Such program instructions can define and/or be part of a discrete software application. In some instances, the remote master control system 300 can execute program instructions in response to receiving an input, such as from the communication interface 306, the user interface 308, or the operations and environment analysis module 310. The data storage unit 304 may also store other information, such as those types described in this disclosure.

In some examples, the data storage unit 304 may serve as storage for information obtained from one or more external sources. For example, data storage unit 304 may store information obtained from one or more of the traditional datacenters 260, a generation station 202, a system associated with the grid, and flexible datacenters 220. As examples only, data storage 304 may include, in whole or in part, local storage, dedicated server-managed storage, network attached storage, and/or cloud-based storage, and/or combinations thereof.

The communication interface 306 can allow the remote master control system 300 to connect to and/or communicate with another component according to one or more protocols. For instance, the communication interface 306 may be used to obtain information related to current, future, and past prices for power, power availability, current and predicted

26

weather conditions, and information regarding the different datacenters (e.g., current workloads at datacenters, types of computing systems available within datacenters, price to obtain power at each datacenter, levels of power storage available and accessible at each datacenter, etc.). In an example, the communication interface 306 can include a wired interface, such as an Ethernet interface or a high-definition serial-digital-interface (HD-SDI). In another example, the communication interface 406 can include a wireless interface, such as a cellular, satellite, WiMAX, or WI-FI interface. A connection can be a direct connection or an indirect connection, the latter being a connection that passes through and/or traverses one or more components, such as such as a router, switcher, or other network device. Likewise, a wireless transmission can be a direct transmission or an indirect transmission. The communication interface 306 may also utilize other types of wireless communication to enable communication with datacenters positioned at various locations.

The communication interface 306 may enable the remote master control system 300 to communicate with the components of the arrangement of FIG. 2. In addition, the communication interface 306 may also be used to communicate with the various datacenters, power sources, and different enterprises submitting computational operations for the datacenters to support.

The user interface 308 can facilitate interaction between the remote master control system 300 and an administrator or user, if applicable. As such, the user interface 308 can include input components such as a keyboard, a keypad, a mouse, a touch-sensitive panel, a microphone, and/or a camera, and/or output components such as a display device (which, for example, can be combined with a touch-sensitive panel), a sound speaker, and/or a haptic feedback system. More generally, the user interface 308 can include hardware and/or software components that facilitate interaction between remote master control system 300 and the user of the system.

In some examples, the user interface 308 may enable the manual examination and/or manipulation of components within the arrangement of FIG. 2. For instance, an administrator or user may use the user interface 308 to check the status of, or change, one or more computational operations, the performance or power consumption at one or more datacenters, the number of tasks remaining within the queue system 312, and other operations. As such, the user interface 308 may provide remote connectivity to one or more systems within the arrangement of FIG. 2.

The operations and environment analysis module 310 represents a component of the remote master control system 300 associated with obtaining and analyzing information to develop instructions/directives for components within the arrangement of FIG. 2. The information analyzed by the operations and environment analysis module 310 can vary within examples and may include the information described above with respect predicting and/or directing the use of BTM power. For instance, the operations and environment analysis module 310 may obtain and access information related to the current power state of computing systems operating as part of the flexible datacenters 220 and other datacenters that the remote master control system 300 has access to. This information may be used to determine when to adjust power usage or mode of one or more computing systems. In addition, the remote master control system 300 may provide instructions a flexible datacenter 220 to cause a subset of the computing systems to transition into a low power mode to consume less power while still performing

US 10,608,433 B1

27

operations at a slower rate. The remote master control system 300 may also use power state information to cause a set of computing systems at a flexible datacenter 220 to operate at a higher power consumption mode. In addition, the remote master control system 300 may transition computing systems into sleep states or power on/off based on information analyzed by the operations and environment analysis module 310.

In some examples, the operations and environment analysis module 310 may use location, weather, activity levels at the flexible datacenters or the generation station, and power cost information to determine control strategies for one or more components in the arrangement of FIG. 2. For instance, the remote master control system 300 may use location information for one or more datacenters to anticipate potential weather conditions that could impact access to power. In addition, the operations and environment analysis module 310 may assist the remote master control system 300 determine whether to transfer computational operations between datacenters based on various economic and power factors.

The queue system 312 represents a queue capable of organizing computational operations to be performed by one or more datacenters. Upon receiving a request to perform a computational operation, the remote master control system 300 may assign the computational operation to the queue until one or more computing systems are available to support the computational operation. The queue system 312 may be used for organizing and transferring computational tasks in real time.

The organizational design of the queue system 312 may vary within examples. In some examples, the queue system 312 may organize indications (e.g., tags, pointers) to sets of computational operations requested by various enterprises. The queue system 312 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 312 may include one or more queues that operate using the FIFO data structure.

In some examples, one or more queues within the queue system 312 may use other designs of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others. Examples using a queue system are further described below with respect to FIG. 9.

In some examples, the remote master control system 300 may be configured to monitor one or more auctions to obtain computational operations for datacenters to support. Particularly, the remote master control system 300 may use resource availability and power prices to develop and submit bids to an external or internal auction system for the right to support particular computational operations. As a result, the remote master control system 300 may identify computational operations that could be supported at one or more flexible datacenters 220 at low costs.

FIG. 4 is a block diagram of a generation station 400, according to one or more example embodiments. Generation station 202 may take the form of generation station 400, or may include less than all components in generation station 400, different components than in generation station 400, and/or more components than in generation station 400. The generation station 400 includes power generation equipment

28

401, a communication interface 408, a behind-the-meter interface 406, a grid interface 404, a user interface 410, a generation station control system 414, and power transformation equipment 402. The power generation equipment 210 may take the form of power generation equipment 401, or may include less than all components in power generation equipment 401, different components than in power generation equipment 401, and/or more components than in power generation equipment 401. Generation station control system 216 may take the form of generation station control system 414, or may include less than all components in generation station control system 414, different components than in generation station control system 414, and/or more components than in generation station control system 414. Some or all of the components generation station 400 may be connected via a communication interface 516. These components are illustrated in FIG. 4 to convey an example configuration for the generation station 400 (corresponding to generation station 202 shown in FIG. 2). In other examples, the generation station 400 may include more or fewer components in other arrangements.

The generation station 400 can correspond to any type of grid-connected utility-scale power producer capable of supplying power to one or more loads. The size, amount of power generated, and other characteristics of the generation station 400 may differ within examples. For instance, the generation station 400 may be a power producer that provides power intermittently. The power generation may depend on monitored power conditions, such as weather at the location of the generation station 400 and other possible conditions. As such, the generation station 400 may be a temporary arrangement, or a permanent facility, configured to supply power. The generation station 400 may supply BTM power to one or more loads and supply metered power to the electrical grid. Particularly, the generation station 400 may supply power to the grid as shown in the arrangement of FIG. 2.

The power generation equipment 401 represents the component or components configured to generate utility-scale power. As such, the power generation equipment 401 may depend on the type of facility that the generation station 400 corresponds to. For instance, the power generation equipment 401 may correspond to electric generators that transform kinetic energy into electricity. The power generation equipment 401 may use electromagnetic induction to generate power. In other examples, the power generation equipment 401 may utilize electrochemistry to transform chemical energy into power. The power generation equipment 401 may use the photovoltaic effect to transform light into electrical energy. In some examples, the power generation equipment 401 may use turbines to generate power. The turbines may be driven by, for example, wind, water, steam or burning gas. Other examples of power production are possible.

The communication interface 408 can enable the generation station 400 to communicate with other components within the arrangement of FIG. 2. As such, the communication interface 408 may operate similarly to the communication interface 306 of the remote master control system 300 and the communication interface 503 of the flexible datacenter 500.

The generation station control system 414 may be one or more computing systems configured to control various aspects of the generation station 400.

The BTM interface 406 is a module configured to enable the power generation equipment 401 to supply BTM power to one or more loads and may include multiple components.

The arrangement of the BTM interface 406 may differ within examples based on various factors, such as the number of flexible datacenters 220 (or 500) coupled to the generation station 400, the proximity of the flexible datacenters 220 (or 500), and the type of generation station 400, among others. In some examples, the BTM interface 406 may be configured to enable power delivery to one or more flexible datacenters positioned near the generation station 400. Alternatively, the BTM interface 406 may also be configured to enable power delivery to one or more flexible datacenters 220 (or 500) positioned remotely from the generation station 400.

The grid interface 404 is a module configured to enable the power generation equipment 401 to supply power to the grid and may include multiple components. As such, the grid interface 404 may couple to one or more transmission lines (e.g., transmission lines 404a shown in FIG. 2) to enable delivery of power to the grid.

The user interface 410 represents an interface that enables administrators and/or other entities to communicate with the generation station 400. As such, the user interface 410 may have a configuration that resembles the configuration of the user interface 308 shown in FIG. 3. An operator may utilize the user interface 410 to control or monitor operations at the generation station 400.

The power transformation equipment 402 represents equipment that can be utilized to enable power delivery from the power generation equipment 401 to the loads and to transmission lines linked to the grid. Example power transformation equipment 402 includes, but is not limited to, transformers, inverters, phase converters, and power conditioners.

FIG. 5 shows a block diagram of a flexible datacenter 500, according to one or more example embodiments. Flexible datacenters 220 may take the form of flexible datacenter 500, or may include less than all components in flexible datacenter 500, different components than in flexible datacenter 500, and/or more components than in flexible datacenter 500. In the example embodiment shown in FIG. 5, the flexible datacenter 500 includes a power input system 502, a communication interface 503, a datacenter control system 504, a power distribution system 506, a climate control system 508, one or more sets of computing systems 512, and a queue system 514. These components are shown connected by a communication bus 528. In other embodiments, the configuration of flexible datacenter 500 can differ, including more or fewer components. In addition, the components within flexible datacenter 500 may be combined or further divided into additional components within other embodiments.

The example configuration shown in FIG. 5 represents one possible configuration for a flexible datacenter. As such, each flexible datacenter may have a different configuration when implemented based on a variety of factors that may influence its design, such as location and temperature that the location, particular uses for the flexible datacenter, source of power supplying computing systems within the flexible datacenter, design influence from an entity (or entities) that implements the flexible datacenter, and space available for the flexible datacenter. Thus, the embodiment of flexible datacenter 220 shown in FIG. 2 represents one possible configuration for a flexible datacenter out of many other possible configurations.

The flexible datacenter 500 may include a design that allows for temporary and/or rapid deployment, setup, and start time for supporting computational operations. For instance, the flexible datacenter 500 may be rapidly

deployed at a location near a source of generation station power (e.g., near a wind farm or solar farm). Rapid deployment may involve positioning the flexible datacenter 500 at a target location and installing and/or configuring one or more racks of computing systems within. The racks may include wheels to enable swift movement of the computing systems. Although the flexible datacenter 500 could theoretically be placed anywhere, transmission losses may be minimized by locating it proximate to BTM power generation.

The physical construction and layout of the flexible datacenter 500 can vary. In some instances, the flexible datacenter 500 may utilize a metal container (e.g., a metal container 602 shown in FIG. 6A). In general, the flexible datacenter 500 may utilize some form of secure weather-proof housing designed to protect interior components from wind, weather, and intrusion. The physical construction and layout of example flexible datacenters are further described with respect to FIGS. 6A-6B.

Within the flexible datacenter 500, various internal components enable the flexible datacenter 500 to utilize power to perform some form of operations. The power input system 502 is a module of the flexible datacenter 500 configured to receive external power and input the power to the different components via assistance from the power distribution system 506. As discussed with respect to FIG. 2, the sources of external power feeding a flexible datacenter can vary in both quantity and type (e.g., the generation stations 202, 400, grid-power, energy storage systems). Power input system 502 includes a BTM power input sub-system 522, and may additionally include other power input sub-systems (e.g., a grid-power input sub-system 524 and/or an energy storage input sub-system 526). In some instances, the quantity of power input sub-systems may depend on the size of the flexible datacenter and the number and/or type of computing systems being powered. In an example embodiment, the flexible datacenter may use grid power as the primary power supply.

In some embodiments, the power input system 502 may include some or all of flexible datacenter Power Equipment 220B. The power input system 502 may be designed to obtain power in different forms (e.g., single phase or three-phase behind-the-meter alternating current ("AC") voltage, and/or direct current ("DC") voltage). As shown, the power input system 502 includes a BTM power input sub-system 522, a grid power input sub-system 524, and an energy input sub-system 526. These sub-systems are included to illustrate example power input sub-systems that the flexible datacenter 500 may utilize, but other examples are possible. In addition, in some instances, these sub-systems may be used simultaneously to supply power to components of the flexible datacenter 500. The sub-systems may also be used based on available power sources.

In some implementations, the BTM power input sub-system 522 may include one or more AC-to-AC step-down transformers used to step down supplied medium-voltage AC to low voltage AC (e.g., 120V to 600V nominal) used to power computing systems 512 and/or other components of flexible datacenter 500. The power input system 502 may also directly receive single-phase low voltage AC from a generation station as BTM power, from grid power, or from a stored energy system such as energy storage system 218. In some implementations, the power input system 502 may provide single-phase AC voltage to the datacenter control system 504 (and/or other components of flexible datacenter 500) independent of power supplied to computing systems 512 to enable the datacenter control system 504 to perform

US 10,608,433 B1

31

management operations for the flexible datacenter 500. For instance, the grid power input sub-system 524 may use grid power to supply power to the datacenter control system 504 to ensure that the datacenter control system 504 can perform control operations and communicate with the remote master control system 300 (or 262) during situations when BTM power is not available. As such, the datacenter control system 504 may utilize power received from the power input system 502 to remain powered to control the operation of flexible datacenter 500, even if the computational operations performed by the computing system 512 are powered intermittently. In some instances, the datacenter control system 504 may switch into a lower power mode to utilize less power while still maintaining the ability to perform some functions.

The power distribution system 506 may distribute incoming power to the various components of the flexible datacenter 500. For instance, the power distribution system 506 may direct power (e.g., single-phase or three-phase AC) to one or more components within flexible datacenter 500. In some embodiments, the power distribution system 506 may include some or all of flexible datacenter Power Equipment 220B.

In some examples, the power input system 502 may provide three phases of three-phase AC voltage to the power distribution system 506. The power distribution system 506 may controllably provide a single phase of AC voltage to each computing system or groups of computing systems 512 disposed within the flexible datacenter 500. The datacenter control system 504 may controllably select which phase of three-phase nominal AC voltage that power distribution system 506 provides to each computing system 512 or groups of computing systems 512. This is one example manner in which the datacenter control system 504 may modulate power delivery (and load at the flexible datacenter 500) by ramping-up flexible datacenter 500 to fully operational status, ramping-down flexible datacenter 500 to offline status (where only datacenter control system 504 remains powered), reducing load by withdrawing power delivery from, or reducing power to, one or more of the computing systems 512 or groups of the computing systems 512, or modulating power factor correction for the generation station 300 (or 202) by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more of the computing systems 512 or groups of the computing systems 512. The datacenter control system 504 may direct power to certain sets of computing systems based on computational operations waiting for computational resources within the queue system 514. In some embodiments, the flexible datacenter 500 may receive BTM DC power to power the computing systems 512.

One of ordinary skill in the art will recognize that a voltage level of three-phase AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase AC voltage, single-phase AC voltage, and DC voltage may vary based on the application or design in accordance with one or more embodiments.

As discussed above, the datacenter control system 504 may perform operations described herein, such as dynamically modulating power delivery to one or more of the computing systems 512 disposed within flexible datacenter 500. For instance, the datacenter control system 504 may modulate power delivery to one or more of the computing

32

systems 512 based on various factors, such as BTM power availability or an operational directive from a generation station 262 or 300 control system, a remote master control system 262 or 300, or a grid operator. In some examples, the datacenter control system 504 may provide computational operations to sets of computing systems 512 and modulate power delivery based on priorities assigned to the computational operations. For instance, an important computational operation (e.g., based on a deadline for execution and/or price paid by an entity) may be assigned to a particular computing system or set of computing systems 512 that has the capacity, computational abilities to support the computational operation. In addition, the datacenter control system 504 may also prioritize power delivery to the computing system or set of computing systems 512.

In some example, the datacenter control system 504 may further provide directives to one or more computing systems to change operations in some manner. For instance, the datacenter control system 504 may cause one or more computing systems 512 to operate at a lower or higher frequency, change clock cycles, or operate in a different power consumption mode (e.g., a low power mode). These abilities may vary depending on types of computing systems 512 available at the flexible datacenter 500. As a result, the datacenter control system 504 may be configured to analyze the computing systems 512 available either on a periodic basis (e.g., during initial set up of the flexible datacenter 500) or in another manner (e.g., when a new computational operation is assigned to the flexible datacenter 500).

The datacenter control system 504 may also implement directives received from the remote master control system 262 or 300. For instance, the remote master control system 262 or 300 may direct the flexible datacenter 500 to switch into a low power mode. As a result, one or more of the computing systems 512 and other components may switch to the low power mode in response.

The datacenter control system 504 may utilize the communication interface 503 to communicate with the remote master control system 262 or 300, other datacenter control systems of other datacenters, and other entities. As such, the communication interface 503 may include components and operate similar to the communication interface 306 of the remote master control system 300 described with respect to FIG. 4.

The flexible datacenter 500 may also include a climate control system 508 to maintain computing systems 512 within a desired operational temperature range. The climate control system 508 may include various components, such as one or more air intake components, an evaporative cooling system, one or more fans, an immersive cooling system, an air conditioning or refrigerant cooling system, and one or more air outtake components. One of ordinary skill in the art will recognize that any suitable heat extraction system configured to maintain the operation of computing systems 512 within the desired operational temperature range may be used.

The flexible datacenter 500 may further include an energy storage system 510. The energy storage system 510 may store energy for subsequent use by computing systems 512 and other components of flexible datacenter 500. For instance, the energy storage system 510 may include a battery system. The battery system may be configured to convert AC voltage to DC voltage and store power in one or more storage cells. In some instances, the battery system may include a DC-to-AC inverter configured to convert DC

US 10,608,433 B1

33

voltage to AC voltage, and may further include an AC phase-converter, to provide AC voltage for use by flexible datacenter 500.

The energy storage system 510 may be configured to serve as a backup source of power for the flexible datacenter 500. For instance, the energy storage system 510 may receive and retain power from a BTM power source at a low cost (or no cost at all). This low-cost power can then be used by the flexible datacenter 500 at a subsequent point, such as when BTM power costs more. Similarly, the energy storage system 510 may also store energy from other sources (e.g., grid power). As such, the energy storage system 510 may be configured to use one or more of the sub-systems of the power input system 502.

In some examples, the energy storage system 510 may be external to the flexible datacenter 500. For instance, the energy storage system 510 may be an external source that multiple flexible datacenters utilize for back-up power.

The computing systems 512 represent various types of computing systems configured to perform computational operations. Performance of computational operations include a variety of tasks that one or more computing systems may perform, such as data storage, calculations, application processing, parallel processing, data manipulation, cryptocurrency mining, and maintenance of a distributed ledger, among others. As shown in FIG. 5, the computing systems 512 may include one or more CPUs 516, one or more GPUs 518, and/or one or more Application-Specific Integrated Circuits (ASIC's) 520. Each type of computing system 512 may be configured to perform particular operations or types of operations.

Due to different performance features and abilities associated with the different types of computing systems, the datacenter control system 504 may determine, maintain, and/or relay this information about the types and/or abilities of the computing systems, quantity of each type, and availability to the remote master control system 262 or 300 on a routine basis (e.g., periodically or on-demand). This way, the remote master control system 262 or 300 may have current information about the abilities of the computing systems 512 when distributing computational operations for performance at one or more flexible datacenters. Particularly, the remote master control system 262 or 300 may assign computational operations based on various factors, such as the types of computing systems available and the type of computing systems required by each computing operation, the availability of the computing systems, whether computing systems can operate in a low power mode, and/or power consumption and/or costs associated with operating the computing systems, among others.

The quantity and arrangement of these computing systems 512 may vary within examples. In some examples, the configuration and quantity of computing systems 512 may depend on various factors, such as the computational tasks that are performed by the flexible datacenter 500. In other examples, the computing systems 512 may include other types of computing systems as well, such as DSPs, SIMDs, neural processors, and/or quantum processors.

As indicated above, the computing systems 512 can perform various computational operations, including in different configurations. For instance, each computing system may perform a particular computational operation unrelated to the operations performed at other computing systems. Groups of the computing systems 512 may also be used to work together to perform computational operations.

In some examples, multiple computing systems may perform the same computational operation in a redundant

34

configuration. This redundant configuration creates a backup that prevents losing progress on the computational operation in situations of a computing failure or intermittent operation of one or more computing systems. In addition, the computing systems 512 may also perform computational operations using a check point system. The check point system may enable a first computing system to perform operations up to a certain point (e.g., a checkpoint) and switch to a second computing system to continue performing the operations from that certain point. The check point system may also enable the datacenter control system 504 to communicate statuses of computational operations to the remote master control system 262 or 300. This can further enable the remote master control system 262 or 300 to transfer computational operations between different flexible datacenters allowing computing systems at the different flexible datacenters to resume support of computational operations based on the check points.

The queue system 514 may operate similar to the queue system 312 of the remote master control system 300 shown in FIG. 3. Particularly, the queue system 514 may help store and organize computational tasks assigned for performance at the flexible datacenter 500. In some examples, the queue system 514 may be part of a distributed queue system such that each flexible datacenter in a fleet of flexible datacenter includes a queue, and each queue system 514 may be able to communicate with other queue systems. In addition, the remote master control system 262 or 300 may be configured to assign computational tasks to the queues located at each flexible datacenter (e.g., the queue system 514 of the flexible datacenter 500). As such, communication between the remote master control system 262 or 300 and the datacenter control system 504 and/or the queue system 514 may allow organization of computational operations for the flexible datacenter 500 to support.

FIG. 6A shows another structural arrangement for a flexible datacenter, according to one or more example embodiments. The particular structural arrangement shown in FIG. 6A may be implemented at flexible datacenter 500. The illustration depicts the flexible datacenter 500 as a mobile container 702 equipped with the power input system 502, the power distribution system 506, the climate control system 508, the datacenter control system 504, and the computing systems 512 arranged on one or more racks 604. These components of flexible datacenter 500 may be arranged and organized according to an example structural region arrangement. As such, the example illustration represents one possible configuration for the flexible datacenter 500, but others are possible within examples.

As discussed above, the structural arrangement of the flexible datacenter 500 may depend on various factors, such as the ability to maintain temperature within the mobile container 602 within a desired temperature range. The desired temperature range may depend on the geographical location of the mobile container 602 and the type and quantity of the computing systems 512 operating within the flexible datacenter 500 as well as other possible factors. As such, the different design elements of the mobile container 602 including the inner contents and positioning of components may depend on factors that aim to maximize the use of space within mobile container 602, lower the amount of power required to cool the computing systems 512, and make setup of the flexible datacenter 500 efficient. For instance, a first flexible datacenter positioned in a cooler geographic region may include less cooling equipment than a second flexible datacenter positioned in a warmer geographic region.

US 10,608,433 B1

35

As shown in FIG. 6A, the mobile container 602 may be a storage trailer disposed on permanent or removable wheels and configured for rapid deployment. In other embodiments, the mobile container 602 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, the mobile container 602 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile flexible datacenter. As such, the flexible datacenter 500 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. And in still other embodiments, the flexible datacenter 500 might not include a mobile container. For example, the flexible datacenter 500 may be situated within a building or another type of stationary environment.

FIG. 6B shows the computing systems 512 in a straight-line configuration for installation within the flexible datacenter 500, according to one or more example embodiments. As indicated above, the flexible datacenter 500 may include a plurality of racks 604, each of which may include one or more computing systems 512 disposed therein. As discussed above, the power input system 502 may provide three phases of AC voltage to the power distribution system 506. In some examples, the power distribution system 506 may controllably provide a single phase of AC voltage to each computing system 512 or group of computing systems 512 disposed within the flexible datacenter 500. As shown in FIG. 6B, for purposes of illustration only, eighteen total racks 604 are divided into a first group of six racks 606, a second group of six racks 608, and a third group of six racks 610, where each rack contains eighteen computing systems 512. The power distribution system (506 of FIG. 5) may, for example, provide a first phase of three-phase AC voltage to the first group of six racks 606, a second phase of three-phase AC voltage to the second group of six racks 608, and a third phase of three-phase AC voltage to the third group of six racks 610. In other embodiments, the quantity of racks and computing systems can vary.

FIG. 7 shows a control distribution system 700 of the flexible datacenter 500 according to one or more example embodiments. The system 700 includes a grid operator 702, a generation station control system 216, a remote master control system 300, and a flexible datacenter 500. As such, the system 700 represents one example configuration for controlling operations of the flexible datacenter 500, but other configurations may include more or fewer components in other arrangements.

The datacenter control system 504 may independently or cooperatively with one or more of the generation station control system 414, the remote master control system 300, and/or the grid operator 702 modulate power at the flexible datacenter 500. During operations, the power delivery to the flexible datacenter 500 may be dynamically adjusted based on conditions or operational directives. The conditions may correspond to economic conditions (e.g., cost for power, aspects of computational operations to be performed), power-related conditions (e.g., availability of the power, the sources offering power), demand response, and/or weather-related conditions, among others.

The generation station control system 414 may be one or more computing systems configured to control various aspects of a generation station (not independently illustrated, e.g., 216 or 400). As such, the generation station control system 414 may communicate with the remote master

36

control system 300 over a networked connection 706 and with the datacenter control system 704 over a networked or other data connection 708.

As discussed with respect to FIGS. 2 and 3, the remote master control system 300 can be one or more computing systems located offsite, but connected via a network connection 710 to the datacenter control system 504. The remote master control system 300 may provide supervisory controls or override control of the flexible datacenter 500 or a fleet of flexible datacenters (not shown).

The grid operator 702 may be one or more computing systems that are configured to control various aspects of the power grid (not independently illustrated) that receives power from the generation station. The grid operator 702 may communicate with the generation station control system 300 over a networked or other data connection 712.

The datacenter control system 504 may monitor BTM power conditions at the generation station and determine when a datacenter ramp-up condition is met. The BTM power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost to obtain power), low priced power, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs. For example, a datacenter control system may analyze future workload and near term weather conditions at the flexible datacenter.

In some instances, the datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702 to go offline or reduce power. As such, the datacenter control system 504 may enable the power input system 502 to provide power to the power distribution system 506 to power the computing systems 512 or a subset thereof.

The datacenter control system 504 may optionally direct one or more computing systems 512 to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems 512 are configured to perform blockchain hashing operations, the datacenter control system 504 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems 512 may be configured to perform high-throughput computing operations and/or high performance computing operations.

The remote master control system 300 may specify to the datacenter control system 504 what sufficient behind-the-meter power availability constitutes, or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 500. In such circumstances, the datacenter control system 504 may provide power to only a subset of computing systems, or operate the plurality of computing systems in a

US 10,608,433 B1

37

lower power mode, that is within the sufficient, but less than full, range of power that is available. In addition, the computing systems 512 may adjust operational frequency, such as performing more or less processes during a given duration. The computing systems 512 may also adjust internal clocks via over-clocking or under-clocking when performing operations.

While the flexible datacenter 500 is online and operational, a datacenter ramp-down condition may be met when there is insufficient or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702. The datacenter control system 504 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by the remote master control system 300 or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently.

An operational directive may be based on current dispatch-ability, forward looking forecasts for when behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the generation station control system 414, the remote master control system 300, or the grid operator 702. For example, the generation station control system 414, the remote master control system 300, or the grid operator 702 may issue an operational directive to flexible datacenter 500 to go offline and power down. When the datacenter ramp-down condition is met, the datacenter control system 504 may disable power delivery to the plurality of computing systems (e.g., 512). The datacenter control system 504 may disable 714 the power input system 502 from providing power (e.g., three-phase nominal AC voltage) to the power distribution system 506 to power down the computing systems 512 while the datacenter control system 504 remains powered and is capable of returning service to operating mode at the flexible datacenter 500 when behind-the-meter power becomes available again.

While the flexible datacenter 500 is online and operational, changed conditions or an operational directive may cause the datacenter control system 504 to modulate power consumption by the flexible datacenter 500. The datacenter control system 504 may determine, or the generation station control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power the flexible datacenter 500. In such situations, the datacenter control system 504 may take steps to reduce or stop power consumption by the flexible datacenter 500 (other than that required to maintain operation of datacenter control system 504).

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically reduce or withdraw power delivery to one or more computing systems 512 to meet the dictate. The datacenter control system 504 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (e.g., 512) to reduce power consumption. The datacenter control system 504 may dynamically reduce the power consumption of one or more computing

38

systems by reducing their operating frequency or forcing them into a lower power mode through a network directive.

Similarly, the flexible datacenter 500 may ramp up power consumption based on various conditions. For instance, the datacenter control system 504 may determine, or the generation control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in greater power generation, availability, or economic feasibility. In such situations, the datacenter control system 504 may take steps to increase power consumption by the flexible datacenter 500.

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to increase power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically increase power delivery to one or more computing systems 512 (or operations at the computing systems 512) to meet the dictate. For instance, one or more computing systems 512 may transition into a higher power mode, which may involve increasing power consumption and/or operation frequency.

One of ordinary skill in the art will recognize that datacenter control system 504 may be configured to have a number of different configurations, such as a number or type or kind of the computing systems 512 that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available behind-the-meter power. As such, the datacenter control system 504 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 8 shows a control distribution system 800 of a fleet of flexible datacenters according to one or more example embodiments. The control distribution system 800 of the flexible datacenter 500 shown and described with respect to FIG. 7 may be extended to a fleet of flexible datacenters as illustrated in FIG. 8. For example, a first generation station (not independently illustrated), such as a wind farm, may include a first plurality of flexible datacenters 802, which may be collocated or distributed across the generation station. A second generation station (not independently illustrated), such as another wind farm or a solar farm, may include a second plurality of flexible datacenters 804, which may be collocated or distributed across the generation station. One of ordinary skill in the art will recognize that the number of flexible datacenters deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more example embodiments.

The remote master control system 300 may provide directive to datacenter control systems of the fleet of flexible datacenters in a similar manner to that shown and described with respect to FIG. 7, with the added flexibility to make high level decisions with respect to fleet that may be counterintuitive to a given station. The remote master control system 300 may make decisions regarding the issuance of operational directives to a given generation station based on, for example, the status of each generation station where flexible datacenters are deployed, the workload distributed across fleet, and the expected computational demand required for one or both of the expected workload and predicted power availability. In addition, the remote master control system 300 may shift workloads from the first plurality of flexible datacenters 802 to the second plurality of flexible datacenters 804 for any reason, including, for

example, a loss of BTM power availability at one generation station and the availability of BTM power at another generation station. As such, the remote master control system 300 may communicate with the generation station control systems 806A, 806B to obtain information that can be used to organize and distribute computational operations to the fleets of flexible datacenters 802, 804.

FIG. 9 shows a queue distribution arrangement for a traditional datacenter 902 and a flexible datacenter 500, according to one or more example embodiments. The arrangement of FIG. 9 includes a flexible datacenter 500, a traditional datacenter 902, a queue system 312, a set of communication links 916, 918, 920A, 920B, and the remote master control system 300. The arrangement of FIG. 9 represents an example configuration scheme that can be used to distribute computing operations using a queue system 312 between the traditional datacenter 902 and one or more flexible datacenters. In other examples, the arrangement of FIG. 9 may include more or fewer components in other potential configurations. For instance, the arrangement of FIG. 9 may not include the queue system 312 or may include routes that bypass the queue system 312.

The arrangement of FIG. 9 may enable computational operations requested to be performed by entities (e.g., companies). As such, the arrangement of FIG. 9 may use the queue system 312 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 500 and the critical traditional datacenter 902. Particularly, the arrangement of FIG. 9 may use the queue system 312 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations among datacenters. As a result, the use of the queue system 312 may reduce time to complete operations and reduce costs.

In some examples, one or more components, such as the datacenter control system 504, the remote master control system 300, the queue system 312, or the control system 936, may be configured to identify situations that may arise where using the flexible datacenter 500 can reduce costs or increase productivity of the system, as compared to using the traditional datacenter 902 for computational operations. For example, a component within the arrangement of FIG. 9 may identify when using behind-the-meter power to power the computing systems 512 within the flexible datacenter 500 is at a lower cost compared to using the computing systems 934 within the traditional datacenter 902 that are powered by grid power. Additionally, a component in the arrangement of FIG. 9 may be configured to determine situations when offloading computational operations from the traditional datacenter 902 indirectly (i.e., via the queue system 312) or directly (i.e., bypassing the queue system 312) to the flexible datacenter 500 can increase the performance allotted to the computational operations requested by an entity (e.g., reduce the time required to complete time-sensitive computational operations).

In some examples, the datacenter control system 504 may monitor activity of the computing systems 512 within the flexible datacenter 500 and use the respective activity levels to determine when to obtain computational operations from the queue system 312. For instance, the datacenter control system 504 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 512 to perform. The various factors may include power availability at the flexible datacenter 500 (e.g., either stored or from a BTM source), availability of the computing systems 512 (e.g., percentage of computing systems avail-

able), type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter 500, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system, among others. The datacenter control system 504 may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems 512 to perform. In such a configuration, the datacenter control system 504 manages the activity of the flexible datacenter 500, including determining when to acquire new sets of computational operations when capacity among the computing systems 512 permit.

In other examples, a component (e.g., the remote master control system 300) within the system may assign or distribute one or more sets of computational operations organized by the queue system 312 to the flexible datacenter 500. For example, the remote master control system 300 may manage the queue system 312, including the distribution of computational operations organized by the queue system 312 to the flexible datacenter 500 and the traditional datacenter 902. The remote master control system 300 may utilize to information described with respect to the Figures above to determine when to assign computational operations to the flexible datacenter 500.

The traditional datacenter 902 may include a power input system 930, a power distribution system 932, a datacenter control system 936, and a set of computing systems 934. The power input system 930 may be configured to receive power from a power grid and distribute the power to the computing systems 934 via the power distribution system 932. The datacenter control system 936 may monitor activity of the computing systems 934 and obtain computational operations to perform from the queue system 312. The datacenter control system 936 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 934 to perform. A component (e.g., the remote master control system 300) within the arrangement of FIG. 9 may assign or distribute one or more sets of computational operations organized by the queue system 312 to the traditional datacenter 902.

The communication link 916 represents one or more links that may serve to connect the flexible datacenter 500, the traditional datacenter 902, and other components within the system (e.g., the remote master control system 300, the queue system 312—connections not shown). In particular, the communication link 916 may enable direct or indirect communication between the flexible datacenter 500 and the traditional datacenter 902. The type of communication link 916 may depend on the locations of the flexible datacenter 500 and the traditional datacenter 902. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

The queue system 312 represents an abstract data type capable of organizing computational operation requests received from entities. As each request for computational operations are received, the queue system 312 may organize the request in some manner for subsequent distribution to a datacenter. Different types of queues can make up the queue system 312 within embodiments. The queue system 312 may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

US 10,608,433 B1

41

In other examples, the queue system 312 may be distributed consisting of multiple queue sub-systems. In the distributed configuration, the queue system 312 may use multiple queue sub-systems to organize different sets of computational operations. Each queue sub-system may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required for performing each set of computational operations. For instance, a first queue sub-system may organize sets of non-intensive computational operations and a second queue sub-system may organize sets of intensive computational operations. In some examples, the queue system 312 may include queue sub-systems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations. In some examples, the queue system 312 may move computational operations between different computing systems or different datacenters in real-time.

Within the arrangement of FIG. 9, the queue system 312 is shown connected to the remote master control system 300 via the communication link 918. In addition, the queue system 312 is also shown connected to the flexible datacenter via the communication 920A and to the traditional datacenter 902 via the communication link 920B. The communication links 918, 920A, 920B may be similar to the communication link 916 and can be various types of communication links within examples.

The queue system 312 may include a computing system configured to organize and maintain queues within the queue system 312. In another example, one or more other components of the system may maintain and support queues within the queue system 312. For instance, the remote master control system 300 may maintain and support the queue system 312. In other examples, multiple components may maintain and support the queue system 312 in a distributed manner, such as a blockchain configuration.

In some embodiments, the remote master control system 300 may serve as an intermediary that facilitates all communication between flexible datacenter 500 and the traditional datacenter 902. Particularly, the traditional datacenter 902 or the flexible datacenter 500 might need to transmit communications to the remote master control system 300 in order to communicate with the other datacenter. As also shown, the remote master control system 300 may connect to the queue system 312 via the communication link 918. Computational operations may be distributed between the queue system 312 and the remote master control system 300 via the communication link 918. The computational operations may be transferred in real-time and mid-performance from one datacenter to another (e.g., from the traditional datacenter 902 to the flexible datacenter 500). In addition, the remote master control system 300 may manage the queue system 312, including providing resources to support queues within the queue system 312.

As a result, the remote master control system 300 may offload some or all of the computational operations assigned to the traditional datacenter 902 to the flexible datacenter 500. This way, the flexible datacenter 500 can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist traditional datacenter 902. The remote master control system 300 may use the queue system 312 to temporarily store and organize the

42

offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter 500) is available to perform them. The flexible datacenter 500 consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to the traditional datacenter 902. The remote master control system 300 may further communicate with the flexible datacenter 500 via communication link 922 and the traditional datacenter 902 via the communication link 924.

FIG. 10A shows method 1000 of dynamic power consumption at a flexible datacenter using behind-the-meter power according to one or more example embodiments. Other example methods may be used to manipulate the power delivery to one or more flexible datacenters.

In step 1010, the datacenter control system, the remote master control system, or another computing system may monitor behind-the-meter power availability. In some embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1020, the datacenter control system or the remote master control system 300 may determine when a datacenter ramp-up condition is met. In some embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the generation station to go offline or reduce power.

In step 1030, the datacenter control system may enable behind-the-meter power delivery to one or more computing systems. In some instances, the remote master control system may directly enable BTM power delivery to computing systems within the flexible system without instructing the datacenter control system.

In step 1040, once ramped-up, the datacenter control system or the remote master control system may direct one or more computing systems to perform predetermined computational operations. In some embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

While operational, the datacenter control system, the remote master control system, or another computing system may receive an operational directive to modulate power consumption. In some embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system or the remote master control system may dynamically reduce power delivery to one or more computing systems or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system or the remote master control system may dynamically adjust power delivery to one or more computing systems to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system may disable power delivery to one or more computing systems.

FIG. 10B shows method 1050 of dynamic power delivery to a flexible datacenter using behind-the-meter power according to one or more embodiments. In step 1060, the datacenter control system or the remote master control system may monitor behind-the-meter power availability. In

US 10,608,433 B1

43

certain embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1070, the datacenter control system or the remote master control system may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the generation station to go offline or reduce power.

In step 1080, the datacenter control system may disable behind-the-meter power delivery to one or more computing systems. In step 1090, once ramped-down, the datacenter control system remains powered and in communication with the remote master control system so that it may dynamically power the flexible datacenter when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system may dynamically modulate power delivery to one or more computing systems of a flexible datacenter based on behind-the-meter power availability or an operational directive. The flexible datacenter may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter may have a blackout state, where all power consumption, including that of the datacenter control system is halted. However, once the flexible datacenter enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system. Generation station conditions or operational directives may cause flexible datacenter to ramp-up, reduce power consumption, change power factor, or ramp-down.

FIG. 11 illustrates a block diagram of a system for implementing control strategies based on a power option agreement, according to one or more embodiments. The system 1100 represents an example arrangement that includes a control system (e.g., the remote master control system 262), a load (e.g., one or more of the datacenters 1102, 1104, and 1106), and a power entity 1140, which may establish and operate in accordance with a power option agreement. Additional arrangements are possible within examples.

In general, a power option agreement is an agreement between a power entity 1140 associated with the delivery of power to a load (e.g., a grid operator, power generation station, or local control station) and the load (e.g., the datacenters 1102-1106). As part of the power option agreement, the load (e.g., load operator, contracting agent for the load, semi-automated control system associated with the load, and/or automated control system associated with the load) provides the power entity 1140 with the right, but not obligation, to reduce the amount of power delivered (e.g., grid power) to the load up to an agreed amount of power during an agreed upon time interval. In order to provide the power entity 1140 with this option, the load needs to be using at least the amount of power subject to the option (e.g., a minimum power threshold). For instance, the load may agree to use at least 1 MW of grid power at all times during a specified 24-hour time interval to provide the power entity 1140 with the option of being able to reduce the amount of power delivered to the load by any amount up to 1 MW at any point during the specified 24-hour time interval. The load may grant the power entity 1140 with this option in exchange for a monetary consideration (e.g., receive power

44

at a reduced price and/or monetary payment if the option is exercised by the power entity).

The power option agreement may be used by the power entity 1140 to reserve the right to reduce the amount of grid power delivered to the load during a set time frame (e.g., the next 24 hours). For instance, the power entity 1140 may exercise a predefined power option to reduce the amount of grid power delivered to the load during a time when the grid power may be better redirected to other loads coupled to the power grid. As such, the power entity 1140 may exercise power option agreements to balance loads coupled to the power grid. In some embodiments, a power option agreement may also specify other parameters, such as costs associated with different levels of power consumption and/or maximum power thresholds for the load to operate according to.

To illustrate an example, a power option agreement may specify that a load (e.g., the datacenters 1102-1106) is required to use at least 10 MW or more at all times during the next 12 hours. Thus, the minimum power threshold according to the power option agreement is 10 MW and this minimum power threshold extends across the time interval of the next 12 hours. In order to comply with the agreement, the load must subsequently operate using 10 MW or more power at all times during the next 12 hours. This way, the load can accommodate a situation where the power entity 1140 exercises the option. Particularly, exercising the option may trigger the load to reduce the amount of power it consumes by an amount up to 10 MW at any point during the 12 hour interval. By establishing this power option agreement, the power entity 1140 can manipulate the amount of power consumed at the load during the next 12 hours by up to 10 MW if power needs to be redirected to another load or a reduction in power consumption is needed for other reasons.

In the example arrangement of the system 1100 shown in FIG. 11, one or more of the datacenters (e.g., the flexible datacenters 1102, 1104, and the traditional datacenter 1106) may operate as the load that is subject to a power option agreement. As the load that is subject to the power option agreement, the datacenters 1102-1106 may execute control instructions in accordance with power target consumption targets that meet or exceed the minimum power thresholds based on the power option agreement.

As shown in FIG. 11, each datacenter 1102-1106 may include a set of computing systems configured to perform computational operations using power from one or more power sources (e.g., BTM power, grid power, and/or grid power subject to a power option agreement). In particular, the flexible datacenter 1102 includes computing systems 1108 arranged into a first set 1114A, a second set 1114B, and a third set 1114C, the flexible datacenter 1104 includes computing systems 1110 arranged into a first set 1116A, a second set 1116B, and a third set 1118B, and the traditional datacenter 1106 includes computing systems 1112 arranged into a first set 1118A, a second set 1118B, and a third set 1118C. Each set of computing systems may include various types of computing systems that can operate in one or more modes.

The different sets of computing systems as well as the multiple datacenters are included in FIG. 11 for illustration purposes. In particular, the variety of computing systems represent different configurations that a load may take while operating in accordance with a power option agreement, and each configuration (as detailed herein) may include ramping up or down power consumption and transferring and performing computational operations between sets of comput-

US 10,608,433 B1

45

ing systems and/or datacenters. In other examples, the load that is subject to a power option agreement may take on other configurations (e.g., a single datacenter **1102-1106**, and/or a single set of computing systems).

The remote master control system **262** may serve as a control system that can determine performance strategies and provide control instructions to the load (e.g., one or more of the datacenters **1102-1106**). In particular, the remote master control system **262** can monitor conditions in concert with the minimum power thresholds and time intervals (e.g., power option data) set forth in, and/or derived from, one or more power option agreements to determine performance strategies that can enable the load to meet the expectations of the power option agreement(s) while also efficiently using power to accomplish computational operations. In some instances, the remote master control system **262** may also be subject to the power option agreement and may adjust its own power consumption based on the power option agreement (e.g., ramp up or down power consumption based on the defined minimum power thresholds during time intervals).

To establish a power option agreement, the remote master control system **262** (or another computing system) may communicate with the power entity **1140**. For instance, the remote master control system **262** may provide a request (e.g., a signal and/or a bid) to the power entity **1140** and receive the terms of one or more power option agreements, or power option data related to power option agreements (e.g., data such as minimum power thresholds and time intervals, but not all terms contained within a potential power option agreement) in response. In some examples, the remote master control system **262** may evaluate one or more conditions prior to establishing a power option agreement to ensure that the conditions could enable the load (e.g., the datacenters **1102-1106**) to operate in accordance with the power option agreement. For instance, the remote master control system **262** may check the quantity and deadlines associated with computational operations assigned to specific datacenters prior to establishing specific datacenters as a load subject to a power option agreement. In some cases, multiple power option agreements may be established. For example, each datacenter **1102-1106** may be subject to a different power option agreement, which may result in the remote master control system **262** managing the power consumption at each of the datacenters **1102-1106** differently.

Within the system **1100** shown in FIG. **11**, the power entity **1140** may represent any type of power entity associated with the delivery of power to the load that is subject to a power option agreement. For instance, the power entity **1140** may be a local station control system, a grid operator, or a power generation source. As such, the power entity **1140** may establish power option agreements with the loads via communication with the loads and/or the remote master control system **262**. For example, the power entity **1140** may obtain and accept a bid from a load trying to engage in a power option agreement with the power entity **1140**. The power entity **1140** is shown with a power option module **1142**, which may be used to establish power option agreements (e.g., fixed-duration **1144** and/or dynamic **1146**).

Once a power option agreement is established, the remote master control system **262** may obtain power option data from the power entity **1140** (or another source) that specifies the power and time expectations of the power entity **1140**. As shown in FIG. **11**, the power entity **1140** includes a power option module **1142**, which may be used to provide power option data to the remote master control system **262** and/or

46

the datacenters **1102-1106**. In particular, the power option data may specify the minimum power threshold or thresholds associated with one or more time intervals for the load to operate at in accordance with based on the power option agreement. The power option data may also specify other constraints that the load should operate in accordance with.

In some examples, the power option data may also include an indication of a monetary penalty that would be imposed upon the load for failure to operate as agreed upon for the power option agreement. In addition, the power option data may also include an indication of a monetary benefit provided to the load operating at power consumption levels that are in accordance with a power option agreement. For instance, monetary benefits could include reduced prices for power, credits for power, and/or monetary payments. In addition, the power option data may include further constraints upon power use, such as one or more maximum power thresholds and corresponding time intervals for the maximum power thresholds.

In some embodiments, the power entity **1140** may correspond to a qualified scheduling entity (QSE). A QSE may submit bids and offers on behalf of resource entities (REs) or load serving entities (LSEs), such as retail electric providers (REPs). QSEs may submit offers to sell and/or bids to buy power (energy) in the Day-Ahead Market (e.g., the next 24 hours) and the Real-Time Market. As such, the remote master control system **262** or another computing system may communicate with one or more QSEs to engage and control one or more loads in accordance with one or more power option agreements.

In some examples, a power option agreement may take the form of a fixed duration power option agreement **1144**. The fixed duration power option agreement **1144** may specify a set of minimum power thresholds and a set of time intervals in advance for an upcoming fixed duration of time covered by the agreement. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. Examples of such association are provided in FIG. **12**. The fixed duration power option agreement may be established in advance of the time period covered by the set of time intervals to enable the remote master control system **262** to prepare performance strategies for the load (e.g., the datacenter(s)) associated with the power option agreement. Thus, the remote master control system **262** may evaluate the fixed duration power option and other monitored conditions to determine performance strategies for a set of computing systems (e.g., one or more datacenters) during the different intervals that satisfy the minimum power thresholds.

In other examples, a power option agreement may take the form of a dynamic power option agreement **1146**. For a dynamic power option agreement **1146**, minimum power thresholds may be provided to the remote master control system **262** in real-time (or near real-time). For instance, a dynamic power option agreement may specify that the power entity **1140** may provide adjustments to minimum power thresholds and corresponding time intervals in real-time to the remote master control system **262**. For example, a dynamic power option agreement may provide power option data that specifies a minimum power threshold for immediate adjustments (e.g., for the next hour).

In an embodiment, a dynamic power option agreement **1146** may involve repeat communication between the remote master control system **262** and the power entity **1140**. Particularly, the power entity **1140** may provide signals to the remote master control system **262** that request power consumption adjustments to be initiated at one or more

US 10,608,433 B1

47

datacenters by the remote master control system 262 over short time intervals, such as across minutes or seconds. For example, the power entity 1140 may communicate to the remote master control system 262 to ramp power consumption down to a particular level within the next 5 minutes. As a result, the remote master control system 262 may provide instructions to one or more datacenters to ramp down power consumption using a linear ramp over the next 5 minutes to meet the particular level specified by the power entity 1140. The remote master control system 262 may monitor the linear ramp down of power consumption and increase or decrease the rate that the datacenter(s) ramp down power use based on projections and updates received from the power entity 1140. As a result, although the ramp down of power consumption may initially be performed in a linear manner to meet a power target threshold, the remote master control system 262 may adjust the rate of power consumption decrease based on updates from the power entity 1140. For example, 25 percent of the overall power consumption ramp down may occur during a first period (e.g., 4 minutes 30 seconds) of the 5 minutes and the remaining 75 percent of the overall power consumption ramp down may occur during the remaining period of the 5 minutes (e.g., the final 30 seconds). The example percentages are included for illustration purposes and can vary within examples based on various parameters, such as additional communication (e.g., adjustments) provided by the power entity 1140.

In further examples, a power option agreement may operate similarly to both a fixed-duration 1144 and a dynamic power option agreement 1146. Particularly, power option data specifying minimum power thresholds and corresponding time intervals may be provided in advance for the entire fixed-duration of time (e.g., the next 24 hours). Additional power option data may then be subsequently provided enabling the remote master control system 262 to make one or more adjustments to accommodate any changes specified within the additional power option data. For instance, additional power option data may indicate that a power entity exercised its option to deliver less power to the load. As a result, the remote master control system may instruct the load to adjust power consumption based on the power entity reducing the power threshold minimum via exercising the option.

As indicated above, the remote master control system 262 may monitor conditions in addition to the constraints set forth in power option data received from the power entity 1140. Particularly, the remote master control system 262 may monitor and analyze a set of conditions (including the power option data) to determine strategies for assigning, transferring, and otherwise managing computational operations using the one or more datacenters 1102-1106. The determined strategies may enable efficient operation by the datacenters while also ensuring that the datacenters operate at target power consumption levels that meet or exceed the minimum power thresholds set forth within one or more power option agreements.

Example monitored conditions include, but are not limited to, power availability 1120, power prices 1122, computing systems parameters 1124, cryptocurrency prices 1126, computational operation parameters 1128, and weather conditions 1129. Power availability 1120 may include determining power consumption ranges at a set of computing systems and/or at one or more datacenters. In addition, power availability 1120 may also involve determining the source or sources of power available at a datacenter. For instance, the remote master control system 262 may identify the types of power sources (e.g., BTM, grid

48

power, and/or a battery system) that a datacenter has available. Power prices 1122 may involve an analysis of the different costs associated with powering a set of computing systems. For instance, the remote master control system 262 may determine cost of power from the grid without a power option agreement relative to the cost power from the grid under the power option agreement. In addition, the remote master control system 262 may also compare the cost of grid power relative to the cost of BTM power when available at a datacenter. The power prices 1122 may also involve comparing the cost of using power at different datacenters to determine which datacenter may perform computational operations at a lower cost.

Monitoring computing system parameters 1124 may involve determining parameters related to the computing systems at one or more datacenters. For instance, the remote master control system 262 may monitor various parameters of the computing systems at a datacenter, such as the abilities and availability of various computing systems, the status of the queue used to store computational operations awaiting performance by the computing systems. The remote master control system 262 may determine types and operation modes of the computing systems, including which computing systems could operate in different modes (e.g., a higher power or a lower power mode) and/or at different hash rates and/or frequencies. The remote master control system 262 may also estimate when computing systems may complete current computational operations and/or how many computational operations are assigned to computing systems.

Monitoring cryptocurrency prices 1126 may involve monitoring the current price of one or more cryptocurrencies, the hash rate and/or estimated power consumption associated with mining each cryptocurrency, and other factors associated with the cryptocurrencies. The remote master control system 262 may use data related to monitoring cryptocurrency prices 1126 to determine whether using computing systems to mine a cryptocurrency generates more revenue than the cost of power required for performance of the mining operations.

The remote master control system 262 may monitor parameters related to computational operations (e.g., computational operation parameters 1128). For example, the remote master control system 262 may monitor parameters related to the computational operations requiring performance and currently being performed, such quantity of operations, estimated time to complete, cost to perform each computational operation, deadlines and priorities associated with each computational operation. In addition, the remote master control system 262 may analyze computational operations to determine if a particular type of computing system may perform the computational operation better than other types of computing systems.

Monitoring weather conditions 1129 may include monitoring for any potential power generation disruption due to emergencies or other events, and changes in temperatures or weather conditions at power generators or datacenters that could affect power generation. As such, the operations and environment analysis module (or another component) of the remote master control system 262 may be configured to monitor one or more conditions described above.

The performance strategy determined by the remote master control system 262 based on the monitored conditions and/or power option data can include control instructions for the load (e.g., the datacenters and/or one or more sets of computing systems). For instance, a performance strategy can specify operating parameters, such as operating frequen-

cies, power consumption targets, operating modes, power on/off and/or standby states, and other operation aspects for computing systems at a datacenter.

The performance strategy can also involve aspects related to the assignment, transfer, and performance of computational operations at the computing systems. For instance, the performance strategy may specify computational operations to be performed at the computing systems, an order for completing computational operations based on priorities associated with the computational operations, and an identification of which computing systems should perform which computational operations. In some instances, priorities may depend on revenue associated with completing each computational operation and deadlines for each computational operation.

The monitored conditions may enable efficient distribution and performance of computational operations among computing systems at one or more datacenters (e.g., datacenters 1102-1106) in ways that can reduce costs and/or time to perform computational operations, take advantage of availability and abilities of computing systems at the datacenters 1102-1106, and/or take advantage in changes in the cost for power at the datacenters 1102-1106. In addition, the monitored conditions may also involve consideration of the power option data to ensure that the computing systems consume enough power to meet minimum power thresholds set forth in one or more power option agreements.

The various monitored conditions described above as well as other potential conditions may change dynamically and with great frequency. Thus, to enable efficient distribution and performance of the computational operations at the datacenters, the remote master control system 262 may be configured to monitor changes in the various conditions to assist with the efficient management and operations of the computing systems at each datacenter. For instance, the remote master control system 262 may engage in wired or wireless communication 1130 with datacenter control systems (e.g., datacenter control system 504) at each datacenter as well as other sources (e.g., the power entity 1140) to monitor for changes in the conditions.

The remote master control system 262 may analyze the different conditions in real-time to modulate operating attributes of computing systems at one or more of the datacenters. By using the monitored conditions, the remote master control system 262 may increase revenue, decrease costs, and/or increase performance of computational operations via various modifications, such as transferring computational operations between datacenters or sets of computing systems within a datacenter and adjusting performance at one or more sets of computing systems (e.g., switching to a low power mode).

In some examples, the traditional datacenter 1106 may be the load subject to a power option agreement. As such, the remote master control system 262 may factor the power option agreement when determining whether to perform computational operations using the computing systems 1112 at the traditional datacenter 1106 and/or transfer computational operations to the computing systems 1108, 1110 at the flexible datacenters 1102, 1104. For instance, the monitored conditions may indicate that the price of grid power is substantially higher than BTM power. As a result, the remote master control system 262 may transfer a subset of computational operations from the traditional datacenter 1106 to the flexible datacenters 1102, 1104. The traditional datacenter 1106 may still have some computational operations to perform to ensure that the traditional datacenter 1106 is

using enough power to meet the minimum power threshold or thresholds set forth in the power option agreement.

In some examples, the remote master control system 262 may monitor the grid frequency signal received from the power entity 1140. When the frequency of the grid deviates a threshold amount (e.g., 0.036 Hz above or below 60 Hz), the remote master control system 262 may adjust performance strategies at the load. In some cases, the remote master control system 262 may adjust the power consumption at the load, the number of miners (or computing systems) operating at the load, and/or the frequency or hash rate, among other possible changes. The remote master control system may readjust performance strategies at the load in response to receiving additional power option data from the power entity 1140 (e.g., an indication that the frequency of the grid is back to 60 Hz). In addition, the remote master control system 262 may communicate changes in operations at the load to the power entity 1140. This way, the power entity 1140 may obtain confirmation that the load is adjusting in accordance with a power option agreement.

In some embodiments, a power generation source (e.g., the generation station 400 shown in FIG. 4) may enter into a power option agreement with a grid operator, which may provide the grid operator with the option to reduce the amount of power that the power source generator can deliver to the grid during a defined time interval. For instance, a wind generation farm may enter into the power option agreement with the grid operator. In addition, the remote master control system 262 may also enter into a power option agreement with the power generation source (e.g., the wind farm) to provide a load that can receive excess power from the power generation source when the grid operator exercises the option and lowers the amount of power that the power generation source can deliver to the grid. Thus, rather than reducing the amount of power produced, the power generation source could exercise an option in the agreement with remote master control system 262 and redirect excess power to one or more loads (e.g., a set of computing systems) that could ramp up power consumption in response. In such situations, the remote master control system 262 maybe able to use the excess power from the power generation source (e.g., BTM power) to perform operations at one or more loads at a low cost (or no cost at all). In addition, the power generation source may benefit from the power option agreement by directing excess power to the load instead of temporarily halting power production.

In some examples, a power option agreement may depend on parameters associated balancing grid capacity and demand. For instance, power option agreements may incentivize power consumption ramping during periods of peak grid power use.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments. The graph 1200 shows power option data arranged according to power 1204 over time 1202. As shown in FIG. 12, time 1202 increases along the X-axis and minimum power thresholds 1204 increase along the Y-axis of the graph 1200. In the example embodiment shown in FIG. 12, the time 1202 increases up to a full day (e.g., 24 hours) in 4 hour increments and the power is shown in MW increasing in intervals of 5 MW. The 24 duration and example minimum power thresholds can differ in other embodiments. Particularly, these values may depend on the terms set forth within the power option agreement.

The graph line 1206 represents sets of minimum power thresholds 1206A, 1206B, 1206C that are specified by

US 10,608,433 B1

51

power option data based on the power option agreement. As shown, the graph line **1206** extends the entire 24 hour duration, which indicates that the set of time intervals associated with minimum power thresholds add up to 24 hours. In other examples, the power option agreement may not include a minimum power threshold during a portion of the duration.

The graph line **1206** of the graph **1200** is further used to illustrate power consumption levels that one or more loads (e.g., a set of computing systems) operating according to the power option agreement may utilize during the 24 hour duration. Particularly, the power quantities above the graph line **1206** represents power levels that the load(s) may consume from the power grid during the 24 hour duration that would satisfy the requirements (i.e., the minimum power thresholds **1206A-1206C**) set forth by the power option agreement. In particular, the power quantities above the graph line **1206** include any power quantity that meets or exceeds the minimum power threshold at that time. By extension, the power quantities positioned below the graph line **1206** represents the amount of power that the load could be directed to reduce power consumption by per the power option agreement.

To further illustrate, an initial minimum power threshold **1206A** is shown associated with the time interval starting at hour 0 and extending to hour 8. In particular, the minimum power threshold **1206A** is set at 5 MW during this time interval. Thus, based on the power option data shown in FIG. **12**, the loads must be able to operate at a target power consumption level that is equal to or greater than the 5 MW minimum power threshold **1206A** at all times during the time interval extending from hour 0 to hour 8, in order to be able to satisfy the power option if it is exercised for that time interval. Similarly, the power entity could reduce the power consumed by loads by any amount up to 5 MW at any point during the time interval from hour 0 to hour 8 in accordance with the power option agreement. For instance, the power entity could exercise its option at any point during this time interval to reduce the power consumed by the loads by 3 MW as a way to load balance the power grid. In response to the power entity exercising its option, the load may then operate using 3 MW less power and/or another strategy determined by a control system factoring additional conditions (e.g., the price of grid power, the revenue that could be generated from mining a cryptocurrency, and/or parameters associated with computational operations awaiting performance).

As further shown in the graph **1200** illustrated in FIG. **12**, the next minimum power threshold **1206B** is associated with the following time interval, which starts at hour 8 and extends until hour 16. During this time interval (hour 8 to hour 16), the load(s) may consume 10 MW or more power since the minimum power threshold **1206B** is now set at 10 MW as shown on the Y-axis of the graph **1200**. In light of the power option data, a control system may determine and provide a performance strategy to the load (e.g., a set of computing systems) that includes a power consumption target that meets or exceeds the minimum power threshold **1206B** (i.e., 10 MW). The performance strategy may depend on the power option data as well as other possible conditions, such as the price of grid power, the availability of computing systems, and/or the type of computing operations, etc. In addition, the power entity could exercise its option to reduce the amount of power consumed by the load by 10 MW or less as represented by the power levels under the minimum threshold **1206B** that extend during the time interval of hour 8 to hour 16.

52

The last minimum power threshold **1206C** is associated with the time interval that starts at hour 16 and extends until hour 24. Similar to the initial minimum power threshold **1206A** associated with the beginning of the graph line **1206**, the last minimum power threshold **1206** is also set at 5 MW. As such, at any point during this interval (hour 16 to hour 24) the loads may consume 5 MW or more to operate in accordance with the power option agreement. As discussed above, by operating at 5 MW or more, the load enables the power consumed from the power grid to be reduced any amount from zero up to 5 MW during this time interval.

When determining the power consumption strategy for a load, a computing system (e.g., the remote master control system **262**) may consider various conditions in addition to the power option data received based on one or more power option agreements. Particularly, the computing system may consider and weigh different conditions in addition to the power option data to determine power consumption targets and/or other control instructions for a load. The conditions may include, but are not limited to, the price of grid power, the price of alternative power sources (e.g., BTM power, stored energy), the revenue associated with mining for one or more cryptocurrencies, parameters related to the computational operations requiring performance (e.g., priorities, deadlines, status of the queue organizing the operations, and/or revenue associated with completing each computational operation), parameters related to the set of computing systems (e.g., types and availabilities of computing systems), and other conditions (e.g., penalties if a minimum power threshold is not met and/or monetary benefits from operating under a power option agreement). By weighing various conditions, the computing system may efficiently manage the set of computing systems, including enabling performance of computational operations cost effectively and/or ensuring that computing systems operate at target power consumption levels that one or more satisfy power option agreements.

In some examples, the computing system may decrease the amount of power that a set of computing systems consumes from one source and while also increasing the amount of power that the set consumes from another source. For instance, the computing system may determine that the price of power grid power is above a threshold price that makes computational operations relatively expensive to perform using grid power. As a result, the computing system may provide control instructions for the computing systems to consume power grid power that matches a minimum power threshold specified by power option data. This may enable the computing systems to satisfy the power option agreement while also avoiding using pricey grid power beyond the minimum amount required per the power option data. In addition, the computing system may instruct some computing systems to switch to a low power mode or temporarily stop until the price of power from the grid decreases. The computing system may instruct one or more computing systems to operate using power from another source (e.g., BTM power and/or stored energy from a battery system) and/or transfer one or more computational operations to another set of computing systems (e.g., a different datacenter).

When the power option agreement is a fixed duration power option agreement, the computing system may receive an indication of all the minimum power thresholds **1206A-1206C** and an indication of the associated time interval altogether and in advance of the duration associated with the power option agreement. By providing all of the minimum power thresholds **1206A-1206C** and the time intervals in

US 10,608,433 B1

53

advance, the computing system may determine a performance strategy for the load that can extend across the entire duration. Particularly, the computing system may factor the minimum power thresholds and associated time intervals as well as other monitored conditions to determine the performance strategy for the total duration. This can enable the computing system to accept and assign computational operations to computing systems in advance while also using a performance strategy that meets the expectations of a power option agreement.

In some examples, the performance strategy determined by the computing system may include control instructions for the set of computing systems to execute if a power option is exercised. For instance, the performance strategy may specify different power consumption targets for the computing systems that depend on whether a power option is exercised during each time interval.

In some instances, the computing system may modify the performance strategy when one or more conditions change enough to warrant a modification. For instance, the computing system may receive an indication of a change in a minimum power threshold (e.g., a decrease in the minimum power threshold) and determine one or more modifications based on the new minimum power threshold and/or other conditions (e.g., a change in the price of power).

In other examples, the power option agreement may be a dynamic power option agreement. Particularly, the load may be subject to a changing minimum power threshold that can vary during a predefined duration associated with the power option agreement. For example, a dynamic power option agreement may specify that the load is subject to a minimum power threshold that may vary from 0 MW up to 5 MW during the next 24 hours and the particular minimum threshold for each hour may depend on power option data received from the power entity during the prior hour. The dynamic power option agreement may further specify the expected response time from the load. For instance, the power option agreement may indicate that an indication of a new minimum power threshold will be provided an hour prior to the start of the minimum power threshold. The computing system, for example, may receive an indication at hour 7 about the increase in the minimum power threshold starting at hour 8. The indication may (or may not) specify the total time interval associated with a new minimum power threshold. For instance, the indication received by the computing system may specify that the 10 MW minimum power threshold extends from hour 8 until hour 16. In other instances, the power option data may indicate that the computing system should abide by the new minimum power threshold until receiving further power option data indicating a change to another new minimum power threshold.

In some examples, the power option data may arrive at the computing system in an unknown order from the power entity with expectations of swift power consumption adjustments by the load. As a result, the power option agreement may require fast ramping of the load to meet changes. Ramping may involve ramping up or down power consumption as well as ramping operating techniques (e.g., adjusting frequency or operation mode).

In some embodiments, the type of power option power agreement may depend on the delivery and content of power option data provided to the load (or a control system controlling the load). For instance, a computing system may receive minimum power thresholds set across an entire duration associated with a power option agreement in advance when the power option agreement is a fixed-

54

duration power option agreement. In other instances, the computing system may receive power option data dynamically and adjust operations in real-time (or near real-time). For instance, the computing system may receive a series of power option data that each specifies minimum power threshold changes during the duration set forth in the dynamic power option agreement. To illustrate an example, the computing system may receive power option data during hour 1 that specifies the minimum power threshold for hour 2, power option data during hour 2 that specifies the minimum power threshold for hour 3, and so on across the duration of the dynamic power option agreement.

In some examples, the minimum power threshold for a time interval may be zero during the duration of a power option agreement. As such, the load may use any amount of power from the power grid in accordance with the power option agreement, including no power at all during this time interval. When the price for power is high during this time frame, the load may ramp down power usage to zero MW to avoid paying the high price for power while still being in compliance with the power option agreement.

FIG. 13 illustrates a method for implementing control strategies based on a fixed-duration power option agreement, according to one or more embodiments. The method 1300 serves as an example and may include other steps within other embodiments. A control system (e.g., the remote master control system 262) may be configured to perform one or more steps of the method 1300. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At step 1302, the method 1300 involves monitoring a set of conditions. For instance, a computing system (e.g., a control system) may monitor various conditions that could impact the performance of operations at one or more loads, including the power consumption targets at the loads. The set of monitored conditions may include a variety of information obtained from one or more external sources, such as one or more datacenters, databases, power generation stations, or types of sources.

Some example conditions include, but are not limited to, the price of grid power, the price and availability of alternative power options (e.g. BTM power, and/or stored energy), parameters of the load (e.g., ramping abilities, type of computing systems, operation modes, etc.), parameters of tasks to be performed using the power at the load (e.g., types, deadlines, priorities, and/or revenue associated with computational operations), availability of other computing systems and their associated costs, and/or revenue associated with mining a cryptocurrency. The computing system may monitor one or more of these conditions as well as others.

At step 1304, the method 1300 involves receiving power option data based, at least in part, on a power option agreement. As discussed above, the computing system (e.g., a remote master control system) may engage in a power option agreement with a power entity. As a result, the computing system may control a load (e.g., a set of computing systems) in accordance with power thresholds and time intervals received from the power entity based on the power option agreement.

In some examples, the power option data may specify a set of minimum power thresholds and a set of time intervals. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. To illustrate an example, the power option data may specify a first minimum power threshold

US 10,608,433 B1

55

associated with a first time interval and a second minimum power threshold associated with a second time interval, with the second time interval subsequent to the first time interval.

The set of time intervals may add up to the duration represented by the power option agreement. For instance, the total duration of the set of time intervals may correspond to a twenty-four hour period (e.g., the next day). In other examples, the power option agreement may span across a different duration (e.g., 12 hours). In additional embodiments, the power option data may specify other information, such as monetary incentives associated with parameters of the power option agreement and/or one or more maximum power thresholds.

At step **1306**, the method **1300** involves determining a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy may be determined responsive to receiving the power option data. In addition, the performance strategy may include a power consumption target for the set of computing systems for each time interval in the set of time intervals. In some examples, each power consumption target is equal to or greater than the minimum power threshold associated with each time interval.

As an example, the performance strategy may specify a first power consumption target for the set of computing systems for a first time interval such that the first power consumption target is equal to or greater than a first minimum power threshold associated with the first time interval and a second power consumption target for the set for a second time interval in a similar manner (i.e., the second power consumption target is equal to or greater than a second minimum power threshold).

In some examples, the performance strategy may include an sequence for the set of computing systems to follow when performing computational operations. The sequence, for example, may be based on priorities associated with the computational operations. In addition, the performance strategy may include one or more power consumption targets that are greater than the minimum power thresholds when the price of power from the power grid is below a threshold price during the time intervals associated with the minimum power thresholds.

The performance strategy may also involve transferring, delaying, or adjusting one or more computational operations performed at the set of computing systems. In addition, the performance strategy may involve adjusting operations at the computing systems. For instance, one or more computing systems may switch modes (e.g., operate at a higher frequency or switch to a low power mode).

In addition, the performance strategy may also specify power consumption targets for the set of computing systems to use if the power option is exercised during an interval. This way, the computing systems may continue to perform computational operations (or suspend performance) based on the power option being exercised.

At step **1308**, the method **1300** involves providing instructions to the set of computing systems to perform one or more computational operations based on the performance strategy. For example, the set of computing systems may operate according to the performance strategy to ensure that the minimum power thresholds are met during the defined time intervals based on the power option agreement.

Some examples may further involve receiving subsequent power option data based, at least in part, on the power option agreement. The subsequent power option data may specify to decrease one or more minimum power thresholds of the

56

set of power thresholds. Responsive to receiving the subsequent power option data, the performance strategy for the set of computing systems may be modified based on a combination of at least a portion of the subsequent power option data and one or more conditions of the monitored conditions. The modified performance strategy may include one or more reduced power consumption targets for the set of computing systems. The amount of the reduction in a power consumption target may depend linearly with the amount that the corresponding minimum power threshold was reduced by. For instance, when a minimum power threshold for a time interval is reduced from 10 MW to 5 MW, the power consumption target for that time interval may be reduced from 10 MW to 5 MW. Instructions may be provided to the set of computing systems to perform computational operations based on the modified performance strategy.

FIG. 14 illustrates a method for implementing control strategies based on a dynamic power option agreement, according to one or more embodiments. The method **1400** serves as an example and may include other steps within other embodiments. Similar to the method **1400**, a control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1400**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At block **1402**, the method **1400** involves monitoring a set of conditions. Similar to block **1302** of the method **1300**, a computing system may monitor various conditions to determine instructions for controlling a set of computing systems.

At block **1404**, the method **1400** involves receiving first power option data based, at least in part, on a power option agreement while monitoring the set of conditions. The first power option data may specify a first minimum power threshold associated with a first time interval. For example, the first power option data may specify a minimum power threshold of 10 MW for the next hour, which may start in an hour or less.

The power option agreement may correspond to a dynamic power option agreement in some examples. When managing a load with respect to a dynamic power option agreement, a computing system may receive power option data specifying changes in minimum power thresholds that a load (e.g., the set of computing systems) may be designated to use in the near term (e.g., the next hour). For example, the computing system may receive power option data during each hour of the duration specified by a power option agreement that indicates a minimum power threshold for the next hour.

At block **1406**, the method **1400** involves providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition. The first control instructions may be provided responsive to receiving the first power option data.

The first control instructions may include a first power consumption target for the set of computing systems for the first time interval. Particularly, the first power consumption target may be equal to or greater than the first minimum power threshold associated with the first time interval. For example, the first power consumption target may be greater than the first minimum power threshold when a cost of power from the power grid is below a threshold price during the first time interval. In other instances, the first power consumption target may be equal to the first minimum power

US 10,608,433 B1

57

threshold when the cost of power from the power grid is greater than the threshold price.

In some examples, control instructions may specify a sequence for the computing systems to follow when performing computational operations. The sequence may be based on priorities associated with each computational operation.

The first control instructions may be determined based on a combination of the first power option data, the price of power from the power grid, and parameters associated with computational operations to be performed at the set of computing systems.

In some examples, the first control instructions may involve ramping up or down power consumption at the set of computing systems. The power consumption may be ramped up or down based on the first minimum power threshold and one or more other conditions (e.g., the price of power).

At block 1408, the method 1400 involves receiving second power option data based, at least in part, on the power option agreement while monitoring the set of conditions. The computing system may receive the second power option data subsequent to receiving the first power option data. The second power option data may specify a second minimum power threshold associated with a second time interval. For example, the second minimum power threshold may be 7 MW over the duration of the upcoming hour. In other examples, the second minimum power threshold may differ as shown in FIG. 12.

In some instances, the computing system may receive the second power option data during the first time interval such that the second time interval overlaps the first time interval. For instance, the computing system may receive the second power option data to enable real-time adjustments to be made to the power consumed at the set of computing systems.

At block 1410, the method 1400 involves providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power option data and at least one condition. The second control instructions may be provided responsive to receiving the second power option data. The second control instructions may specify a second power consumption target for the set of computing systems for the second time interval. The second power consumption target may be equal to or greater than the second minimum power threshold associated with the second time interval.

In some examples, the computing system may provide a request to a QSE to determine the power option agreement. As such, the computing system may receive power option data (e.g., the first and second power option data) in response to providing the request to the QSE.

The computing system may monitor the price of power from the power grid, and the global mining hash rate and a price for a cryptocurrency (e.g., Bitcoin), among other conditions. The computing system may determine control instructions (e.g., the first and/or second control instructions) based on a combination of power option data, the price of power from the power grid, and the global mining hash rate and the price for the cryptocurrency. For instance, the computing system may cause one or more computing systems (e.g., a subset of computing systems) to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

58

Advantages of one or more embodiments of the present invention may include one or more of the following:

One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive low cost or unutilized power behind-the-meter when it is available.

One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations across multiple datacenters.

One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

One or more embodiments of the present invention allows for management and distribution of computational operations at computing systems across a fleet of datacenters such that the performance of the computational operations take advantages of increased efficiency and decreased costs.

It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

US 10,608,433 B1

59

What is claimed is:

1. A system comprising:
 - a set of computing systems, wherein the set of computing systems is configured to perform computational operations using power from a power grid;
 - a control system configured to:
 - monitor a set of conditions;
 - receive power option data based, at least in part, on a power option agreement, wherein the power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, wherein each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals;
 - responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions, wherein the performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, wherein each power consumption target is equal to or greater than the minimum power threshold associated with each time interval; and
 - provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.
2. The system of claim 1, wherein the control system is configured to monitor the set of conditions comprising:
 - a price of power from the power grid; and
 - a plurality of parameters associated with one or more computational operations to be performed at the set of computing systems.
3. The system of claim 2, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at least the portion option data, the price of power from the power grid, and the plurality of parameters associated with the one or more computational operations.
4. The system of claim 3, wherein the performance strategy further comprises:
 - an order for the set of computing systems to follow when performing the one or more computational operations, wherein the order is based on respective priorities associated with the one or more computational operations.
5. The system of claim 4, wherein the performance strategy further comprises:
 - at least one power consumption target that is greater than a minimum power threshold when the price of power from the power grid is below a threshold price during the time interval associated with the minimum power threshold.
6. The system of claim 1, wherein the control system is further configured to:
 - receive subsequent power option data based, at least in part, on the power option agreement,
 - wherein the subsequent power option data specify to decrease one or more minimum power thresholds of the set of minimum power thresholds.
7. The system of claim 6, wherein the control system is further configured to:
 - responsive to receiving the subsequent power option data,
 - modify the performance strategy for the set of computing systems based on a combination of at least the

60

- portion of the subsequent power option data and at least one condition in the set of conditions,
- wherein the modified performance strategy comprises one or more reduced power consumption targets for the set of computing systems.
8. The system of claim 7, wherein the control system is further configured to:
 - provide instructions to the set of computing systems to perform the one or more computational operations based on the modified performance strategy.
9. The system of claim 1, wherein the control system is a remote master control system positioned remotely from the set of computing systems.
10. The system of claim 1, wherein the control system is a mobile computing device.
11. The system of claim 1, wherein the control system is configured to receive the power option data while monitoring the set of conditions.
12. The system of claim 1, wherein the control system is further configured to:
 - provide a request to a qualified scheduling entity (QSE) to determine the power option agreement; and
 - receive power option data in response to providing the request to the QSE.
13. The system of claim 1, wherein the power option data specify: (i) a first minimum power threshold associated with a first time interval in the set of time intervals, and (ii) a second minimum power threshold associated with a second time interval in the set of time intervals,
 - wherein the second time interval is subsequent to the first time interval.
14. The system of claim 13, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems such that the performance strategy comprises:
 - a first power consumption target for the set of computing systems for the first time interval, wherein the first power consumption target is equal to or greater than the first minimum power threshold; and
 - a second power consumption target for the set of computing systems for the second time interval, wherein the second power consumption target is equal to or greater than the second minimum power threshold.
15. The system of claim 1, wherein a total duration of the set of time intervals corresponds to a twenty-four hour period.
16. The system of claim 1, wherein the set of conditions monitored by the control system further comprise:
 - a price of power from the power grid; and
 - a global mining hash rate and a price for a cryptocurrency; and
 wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at the portion of the power option data, the price of power from the power grid, the global mining hash rate and the price for the cryptocurrency,
 - wherein the performance strategy specifies for at least a subset of the set of computing systems to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

US 10,608,433 B1

61

17. A method comprising:
 monitoring, by a computing system, a set of conditions;
 receiving, at the computing system, power option data
 based, at least in part, on a power option agreement,
 wherein the power option data specify: (i) a set of
 minimum power thresholds, and (ii) a set of time
 intervals, wherein each minimum power threshold in
 the set of minimum power thresholds is associated with
 a time interval in the set of time intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.

18. The method of claim 17, wherein determining the
 performance strategy for the set of computing systems
 comprises:
 identifying information about the set of computing sys-
 tems; and
 determining the performance strategy to further comprise
 instructions for at least a subset of the set of computing
 systems to operate at an increased frequency based on
 a combination of at least the portion of the power
 option data and the information about the set of com-
 puting systems.

19. The method of claim 17, further comprising:
 receiving subsequent power option data based, at least in
 part, on the power option agreement, wherein the
 subsequent power option data specify to decrease one
 or more minimum power thresholds of the set of
 minimum power thresholds;

62

responsive to receiving the subsequent power option data,
 modifying the performance strategy for the set of
 computing systems based on a combination of at least
 the portion of the subsequent power option data and at
 least one condition in the set of conditions, wherein the
 modified performance strategy comprises one or more
 reduced power consumption targets for the set of com-
 puting systems; and
 providing instructions to the set of computing systems to
 perform the one or more computational operations
 based on the modified performance strategy.

20. A non-transitory computer readable medium having
 stored therein instructions executable by one or more pro-
 cessors to cause a computing system to perform functions
 comprising:
 monitoring a set of conditions;
 receiving power option data based, at least in part, on a
 power option agreement, wherein the power option
 data specify: (i) a set of minimum power thresholds,
 and (ii) a set of time intervals, wherein each minimum
 power threshold in the set of minimum power thresh-
 olds is associated with a time interval in the set of time
 intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.

* * * * *

EXHIBIT B

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

**Title of
Invention**

Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration
Power Option Agreement

As the below named inventor, I hereby declare that:

This declaration
is directed to:

☒

The attached application, or

☐

United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001
by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Michael T. McNamara

Date (Optional): Dec 2, 2019

Signature: 

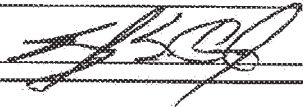
Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration Power Option Agreement
<p>As the below named inventor, I hereby declare that:</p> <p>This declaration is directed to: <input checked="" type="checkbox"/> The attached application, or <input type="checkbox"/> United States application or PCT international application number _____ filed on _____</p> <p>The above-identified application was made or authorized to be made by me.</p> <p>I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.</p> <p>I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.</p> <p style="text-align: center;">WARNING:</p> <p>Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.</p>	
<p>LEGAL NAME OF INVENTOR</p> <p>Inventor: <u>Raymond E. Cline Jr.</u> Date (Optional): <u>12/3/2019</u></p> <p>Signature: </p>	
<p>Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.</p>	

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

~~Appx266~~

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,)	
)	
Plaintiffs,)	
)	
v.)	C.A. No. 21-534-MN
)	
LANCIUM LLC, MICHAEL T.)	JURY TRIAL DEMANDED
MCNAMARA, and RAYMOND E. CLINE,)	
JR.)	
)	
Defendants.)	

AMENDED COMPLAINT

Plaintiffs BearBox LLC (“BearBox”) and Austin Storms (collectively, “Plaintiffs”) bring this action against Lancium LLC (“Lancium”), Michael T. McNamara, and Raymond E. Cline, Jr. (collectively “Defendants”) to correct the inventorship of U.S. Patent No. 10,608,433 (the “’433 Patent”) and to recover damages, injunctive relief, declaratory relief, and other remedies for Defendants’ wrongful actions to obtain, misuse, disclose, and claim as their own Plaintiffs’ proprietary cryptocurrency mining technology. Plaintiffs further allege as follows:

INTRODUCTION

1. This case is about the Defendants’ theft of inventions that rightfully belong to Plaintiffs.
2. Plaintiffs developed proprietary technology relating to cryptocurrency mining systems (the “BearBox Technology”). By way of background, the BearBox Technology generally relates to an energy-efficient cryptocurrency mining system and related methods that reduce the inefficiency and environmental impact of energy-expensive mining operations by better utilizing available energy resources to increase stability of the energy grid, minimize a

mining operation's impact on peak-demand, and also alleviate energy over-supply conditions. The BearBox Technology can be used to mine cryptocurrency, such as Bitcoin.

3. The Defendants induced the Plaintiffs to disclose the BearBox Technology to them under the guise of a possible business deal between Defendants and Plaintiffs to jointly commercialize the BearBox Technology. Before disclosing the BearBox Technology to Defendants, Plaintiffs obtained assurances of confidentiality from Defendants.

4. The Defendants stole the BearBox Technology from Plaintiffs by converting and misappropriating it and claiming it as their own. Defendants filed a U.S. patent application that wrongfully disclosed the BearBox Technology to the U.S. Patent and Trademark Office and ultimately to the public. The claimed subject matter of the '433 Patent falls fully within the scope of the BearBox Technology. And by obtaining the '433 Patent with claims directed to the BearBox Technology, the Defendants have wrongfully obtained a patent covering the BearBox Technology and wrongfully claimed the BearBox Technology as their own.

5. Plaintiffs bring this action to correct the named inventors on the '433 Patent. The inventions claimed in the '433 Patent are inventions conceived by Storms, founder and president of BearBox.

PARTIES

6. Plaintiff BearBox LLC ("BearBox") is a limited liability company organized and existing under the laws of Louisiana with its principal place of business at 4422 Highway 22, Mandeville, Louisiana 70471.

7. Plaintiff Austin Storms is an individual residing in Mandeville, Louisiana.

8. On information and belief, Defendant Lancium is a Delaware limited liability company with its principal place of business at 6006 Thomas Rd, Houston, Texas 77041. On

information and belief, Lancium has a registered agent capable of accepting service in this district, Harvard Business Services, Inc. with a place of business at 16192 Coastal Highway, Lewes, DE 19958.

9. On information and belief, Defendant Michael T. McNamara is the Chief Executive Officer and a founder of Lancium and resides in Newport Beach, California. Defendant McNamara is named as a purported inventor on the face of the '433 Patent.

10. On information and belief, Defendant Raymond E. Cline, Jr. is the Chief Computing Officer of Lancium and resides in Houston, Texas. Defendant Cline is named as a purported inventor on the face of the '433 Patent.

JURISDICTION

11. This is an action seeking correction of the named inventors of a United States patent under 35 U.S.C. § 256. As such, this action arises under the laws of the United States.

12. This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under an Act of Congress relating to patents, specifically 35 U.S.C. § 256.

.

13. The Court has supplemental jurisdiction under 28 U.S.C. § 1367 over all asserted claims under state law because those claims are so related to the claims in this action that arise under federal law that they form part of the same case or controversy.

14. The Court also has jurisdiction pursuant to 28 U.S.C. § 1332, as complete diversity of citizenship exists among the parties, and the amount in controversy exceeds \$75,000. Plaintiff BearBox is a citizen of the State of Louisiana because it is organized under the laws of the State of Louisiana and has its principal place of business in the State of Louisiana. Plaintiff

Storms is a citizen of the State of Louisiana because he resides in the State of Louisiana. In contrast, none of the Defendants are citizens of the State of Louisiana. Defendant Lancium is a citizen of the States of Delaware and Texas because it is organized under the laws of the State of Delaware and has its principal place of business in the State of Texas. Defendant McNamara is a citizen of the State of California because he resides in the State of California. Defendant Cline is a citizen of the State of Texas because he resides in the State of Texas. Therefore, because the Plaintiffs are both citizens of the State of Louisiana (and no other states) for purposes of diversity jurisdiction, and none of the Defendants are citizens of the State of Louisiana, complete diversity exists among the parties.

15. This Court has general personal jurisdiction over Lancium because it is organized under the laws of the State of Delaware and because it maintains an ongoing presence in this District at least through its registered agent.

16. This Court has specific personal jurisdiction over each of Defendants McNamara and Cline at least under Title 6 of the Delaware Code, § 18-109(a).

17. On information and belief, Defendant McNamara is the Chief Executive Officer of Lancium. On information and belief, as the Chief Executive Offer, McNamara participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

18. McNamara is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from the interests of Lancium and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claims against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant

McNamara together. Plaintiffs' claims against Defendant McNamara arise out of his exercise of his powers as Chief Executive Officer of Lancium.

19. On information and belief, Defendant Cline is the Chief Computing Officer of Lancium. On information and belief, as the Chief Computing Officer, Cline participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

20. Cline is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from Lancium's interest and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claim against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant Cline together.

Plaintiffs' claims against Defendant Cline arise out of his exercise of his powers as Chief Computing Officer of Lancium.

21. The actions of Defendants McNamara and Cline establish sufficient minimum contacts with Delaware under Delaware law and the United States Constitution to give this Court personal jurisdiction over each of them.

22. As described below, each Defendant has committed acts giving rise to this action.

VENUE

23. Venue is proper in this District under 28 U.S.C. § 1391(b)(3) because there is no district in which an action may otherwise be brought as provided in § 1391(b) and Defendant Lancium is subject to the Court's personal jurisdiction with respect to this action.

PLAINTIFFS' PROPRIETARY CRYPTOCURRENCY MINING TECHNOLOGY

24. As of 2018, the amount of energy required to process computer algorithms to mine cryptocurrencies like Bitcoin was three times greater than the energy required to physically mine gold. Conventional mining of “copper, gold, platinum, and rare earth oxides are 4, 5, 7, and 9 megajoules to generate one U.S. dollars,” while “it costs an average of 17 megajoules to mine \$1 worth of bitcoin.”¹ The large amount of energy required to mine cryptocurrencies can make such mining financially prohibitive, and even when financially lucrative, the large energy requirements make cryptocurrency mining harmful to the global environment, with studies showing carbon dioxide emissions from cryptocurrency mining “single-handedly rais[ing] global temperatures by 2 degrees by 2023.” *Id.*

25. At the same time, some forms of electrical power generation are terribly inefficient. When producers of electrical power are unable to quickly adjust their operations in response to dynamically changing grid conditions, these producers frequently sell power at low or even negative prices until demand and market prices increase.

26. Because cryptocurrency mining is a computationally demanding process, it requires significant energy. As a result, industrial-scale cryptocurrency mining places a large energy burden on the power grid, driving demand and costs as well as increasing the likelihood of grid component failure.

27. In late 2018 and early 2019, Austin Storms sought to address these problems by developing energy-efficient cryptocurrency mining systems and methods that reduce the environmental impact of energy-intensive mining operations. Storms conceived of a system that

¹ <https://www.marketwatch.com/story/mining-bitcoin-is-3-times-more-expensive-than-mining-gold-research-paper-finds-2018-11-06>

better uses available energy resources to increase the stability of the energy grid, minimize a mining operation's impact on peak-demand, and alleviate energy over supply conditions, all while decreasing the overall energy costs of the mining operation and increasing its profitability.

28. Austin Storms conceived of and developed the BearBox Technology. Storms is the president and founder of BearBox. The BearBox Technology includes hardware and software components. Structurally, the BearBox Technology includes a housing for a plurality of miners (such as ASICs, graphics cards, or the like) under the direction of a smart controller(s).

29. The smart controller monitors various external factors, such as current and expected energy demand and pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the system may determine whether conditions are appropriate to mine cryptocurrency and, if so, subsequently mines the cryptocurrency. Optionally, the system also includes other components for cooling, air-filtration, and related features.

30. In the BearBox Technology, a controller (such as a power distribution unit, network interface, or the like) monitors various external factors, such as current and expected energy demand/pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the controller(s) determines appropriate times to mine cryptocurrency in accordance with a desired performance strategy (for example, profitability thresholds). At the appropriate times, the controller initiates mining, for example, by powering on the miners.

**DEFENDANTS WRONGFULLY CLAIM THE
BEARBOX TECHNOLOGY AS THEIR OWN**

31. In May 2019, Storms attended the Fidelity FCAT Mining Summit in Boston, Massachusetts on behalf of BearBox to promote the BearBox Technology and seek potential customers for his revolutionary system.

32. While at the conference, Storms met Defendant McNamara. Defendant McNamara showed immediate interest in the BearBox Technology. Under the rouse of a potential business relationship, McNamara pumped Storms for details about the BearBox Technology over the course of several exchanges, which included conversations, emails, and text messages about the BearBox Technology. Storms took McNamara to dinner where McNamara continued to pump Storms for details about the BearBox Technology. At all times before and during Storms's disclosure of this information, Storms told McNamara that the BearBox Technology was confidential, and Storms relied on McNamara's good faith assurances that he would keep confidential the information he received from Storms about the BearBox Technology.

33. Following the conference, McNamara continued to press Storms for additional details about the BearBox Technology via text messaging and email. Again relying on Defendant McNamara's assurances of confidentiality, Storms provided annotated system diagrams, component specifications, and modeled data sets to mimic real-world Bitcoin and energy prices. Storms included express confidentiality notices in his communications with Defendant McNamara.

34. After Storms disclosed the BearBox Technology to McNamara, McNamara abruptly ended all communications with Storms.

35. Storms last communicated with McNamara on May 9, 2019 via e-mail, and after sending that message, Storms did not hear from McNamara again.

36. At that time, Storms understood that McNamara was not interested in investing in the BearBox Technology. He had no reason to suspect that McNamara would steal the BearBox Technology and claim it as his own.

37. On information and belief, Defendants filed U.S. provisional patent application No. 62/927,119 on October 28, 2019, naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

38. In addition to falsely claiming to be the inventors of the inventions disclosed in the application, Defendants wrongfully disclosed, without authorization, the confidential BearBox Technology to the United States Patent and Trademark Office.

39. Likewise, on December 4, 2019, Defendants filed U.S. Patent Application Serial No. 16/702,931, once again naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

40. The '433 Patent issued on March 31, 2020 naming Defendants McNamara and Cline as the sole purported inventors on the face of the patent. A true and correct copy of the '433 Patent is attached hereto as Exhibit A.

41. The inventions claimed in the '433 patent fall within the scope of the BearBox Technology, yet Defendants falsely identified themselves as the inventors of the claimed inventions, when, in fact, Storms is the sole inventor of the claimed inventions.

42. On information and belief, McNamara and Cline assigned their purported rights in the '433 patent to Lancium. On information and belief, at all times, Lancium was aware that

McNamara and Cline, both officers of Lancium, were not the rightful inventors of the BearBox Technology disclosed in the patent and the inventions claimed in the patent.

43. Defendants McNamara and Cline each submitted signed declarations falsely swearing that they were “an original joint inventor” of the claimed subject matter . A true and correct copy of Defendant McNamara’s and Defendant Cline’s declarations are attached as Exhibit B.

44. On August 14, 2020, Lancium filed a lawsuit in the U.S. District Court for the Western District of Texas against Layer1 Technologies, Inc. (“Layer1”) asserting that Layer1 infringes the ’433 patent. That case is captioned *Lancium LLC v. Layer1 Technologies, Inc.*, Case No. 6:20-cv-739 (W.D. Texas) (the “Layer1 Lawsuit”).

45. As part of the Layer1 Lawsuit, Defendants falsely asserted that McNamara and Cline are the sole inventors of the inventions claimed in the ’433 patent.

46. Plaintiffs became aware of Defendants’ wrongful use of the BearBox Technology on or about August 17, 2020, when they learned about the Layer1 Lawsuit through a press release dated August 14, 2020, posted by Lancium on PRNewswire. That press release is available at the following URL: <https://www.prnewswire.com/news-releases/controllable-load-resource-clr-market-leader-lancium-files-patent-infringement-lawsuit-against-layer1-301112687.html>.

47. Before seeing the August 14, 2020 press release, Plaintiffs were unaware of Defendants’ wrongful use of the BearBox Technology and was unaware of the ’433 patent.

48. On March 5, 2021, Lancium and Layer1 entered a Stipulation to Dismiss with Prejudice in the Layer1 Lawsuit. According to the stipulation, the parties had entered a Settlement Agreement to resolve the Layer1 Lawsuit.

49. According to a press release issued by Lancium on March 8, 2021, Lancium and Layer 1 “have entered into a mutually beneficial partnership. Layer1 has licensed Lancium’s intellectual property and Lancium will provide Smart Response™ software and services to Layer1.” The press release is available at the following URL: <https://www.prnewswire.com/news-releases/lancium-and-layer1-settle-patent-infringement-suit-301242602.html>

50. On information and belief, as part of the Settlement Agreement between Lancium and Layer1 to settle the Layer1 Lawsuit, Lancium received and continues to receive valuable consideration from Layer1, all of which rightly belongs to Plaintiffs, the rightful owners of the inventions claimed in the ’433 Patent.

COUNT I
CORRECTION OF INVENTORSHIP FOR THE ’433 PATENT:
AUSTIN STORMS AS SOLE INVENTOR

51. Plaintiffs incorporate the above paragraphs by reference.

52. Storms is the sole inventor of the subject matter claimed in the ’433 Patent.

53. Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the ’433 patent and the currently listed inventors on the ’433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms or BearBox.

54. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the ’433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT II
IN THE ALTERNATIVE, CORRECTION OF INVENTORSHIP FOR THE '433
PATENT: AUSTIN STORMS AS JOINT INVENTOR WITH THE CURRENTLY
NAMED INVENTORS

55. Plaintiffs incorporates the above paragraphs by reference.

56. In the alternative, Storms is a joint inventor of the subject matter claimed in the '433 Patent and should be added to the individuals currently named as inventors on the '433 Patent.

57. Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the '433 patent and the currently listed inventors on the '433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms.

58. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the '433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT III
CONVERSION BY LANCIUM, MCNAMARA, AND CLINE

59. Plaintiffs incorporates the above paragraphs by reference.

60. Austin Storms, in his capacity as founder and President of BearBox, conceived, developed, and reduced to practice the BearBox Technology. Plaintiffs own the BearBox Technology, related know-how, and related intellectual property. Plaintiffs owned this property during all relevant time periods in this suit. Information on the BearBox Technology was provided to Defendants solely for the purposes of evaluation for a potential business relationship and under strict confidentiality obligations.

61. Defendants assumed dominion and control over the BearBox Technology by claiming it as their own in the '433 patent. Through their wrongful conduct in obtaining the '433 Patent and claiming the BearBox Technology as their own, the Defendants have wrongfully obtained the purported ability to exclude Plaintiffs and others from using the BearBox Technology. This constitutes unauthorized and unlawful conversion by Defendants.

62. As a result of Defendants' wrongful actions, Plaintiffs will suffer imminent and irreparable damages in an amount to be proven at trial. In particular, Plaintiffs have been damaged by losing valuable intellectual property from which Plaintiffs would have derived substantial revenue via licensing and/or selling patented products.

**COUNT IV
UNJUST ENRICHMENT BY LANCIUM, MCNAMARA, AND CLINE**

63. Plaintiffs incorporate the above paragraphs by reference.

64. Plaintiffs conferred a benefit on Defendants by providing them valuable intellectual property about cryptocurrency mining systems and related confidential information and materials under the boundaries of a potential collaboration between BearBox and Lancium.

65. Defendants accepted that cryptocurrency mining intellectual property and, indeed, continuously asked Storms to provide more information and materials, having recognized the benefit that Defendants received by having access to the BearBox Technology.

66. Defendants accepted and retained the BearBox Technology, and used it to their own advantage, at Plaintiffs' expense.

67. Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct. In particular, Defendants have unlawfully used Plaintiffs' property by asserting inventorship over the BearBox Technology, and deriving an unjust benefit from

exploiting Storms's cryptocurrency mining inventions. It would be inequitable for Defendants to retain these benefits under these circumstances.

68. Plaintiffs have incurred, and continue to incur, detriment in the form of loss of money and property as a result of Defendants' wrongful use of Plaintiffs' intellectual property, including the right to any patent based on their own intellectual property. The intellectual property, including the right to any patents based on Plaintiffs' intellectual property and to any patent documents (including assignment documents), U.S. and foreign, are unique and there is no adequate remedy at law.

69. The harm to Plaintiffs is continuous, substantial, and irreparable.

COUNT V NEGLIGENT MISREPRESENTATION BY LANCIUM AND MCNAMARA

70. Plaintiffs incorporate the above paragraphs by reference.

71. In connection with the potential work involving cryptocurrency mining systems and related methods, Storms told Defendant McNamara that the cryptocurrency mining systems and related methods were proprietary to Plaintiffs and not to be used or shared outside of Lancium. Defendant McNamara gave his word that he would abide by this confidentiality. On information and belief, Defendant McNamara agreed to keep the BearBox Technology confidential despite later recklessly incorporating the BearBox Technology into his own patent applications and swearing, as recently as December 4, 2019, that he is an inventor of the BearBox Technology. Storms relied on Defendant McNamara's assurances of confidentiality and continued to share details about the BearBox Technology with Defendants.

72. If Plaintiffs had known that Defendants would secretly incorporate the BearBox Technology into Defendants' own patent applications to claim them as Defendants' intellectual

property, Plaintiffs would not have continued working with and sharing intellectual property with Defendants.

73. Plaintiffs suffered a pecuniary loss based on this reliance including the loss of potential patent rights, and the costs of Plaintiffs' know-how converted under the guise of a potential business relationship.

JURY DEMAND

74. Under Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiffs respectfully demand a trial by jury on all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, BearBox respectfully requests the following relief:

A. An order that the Director of the United States Patent and Trademark Office correct the inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

B. Alternatively, an order that Defendants sign the requisite documents to correct inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

C. A declaration that Austin Storms is the sole inventor, or, in the alternative, is a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

D. A preliminary and a permanent injunction enjoining Defendants Lancium, McNamara, and Cline from asserting that McNamara or Cline are inventors of the '433 Patent in violation of the United States federal patent laws;

E. An order that Defendants immediately transfer to Plaintiffs all right, title, and interest in all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

F. An order for a constructive trust over all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

G. Financial relief including damages, consequential damages, disgorgement of Defendants' ill-gotten profits, Defendants' unjust enrichment, reasonable royalty damages, lost profits damages, reliance damages, and/or all other appropriate financial relief, all in an amount to be determined at trial, with interest;

H. An award of the amount by which Defendants have been unjustly enriched by their actions set forth in this Complaint and their purported ownership of patents covering Plaintiffs' intellectual property;

I. A finding that this is an exceptional case warranting imposition of attorney fees against Defendants and an award to Plaintiffs of its reasonable costs and attorney fees incurred in bringing this action pursuant to 35 U.S.C. § 285; and

J. An award of such further relief at law or in equity, such as preliminary and/or permanent injunctive relief, as the Court deems just and proper.

ASHBY & GEDDES

/s/ Andrew C. Mayo

Andrew C. Mayo (#5207)
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19899
(302) 654-1888
amayo@ashbygeddes.com

Attorneys for Plaintiffs
BearBox LLC and Austin Storms

Of Counsel:

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati, III
Chelsea M. Murray
MARSHALL, GERSTEIN & BORUN LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
(312) 474-6300

Dated: May 24, 2021

EXHIBIT A

(12) **United States Patent**
McNamara et al.

(10) **Patent No.:** **US 10,608,433 B1**
 (45) **Date of Patent:** **Mar. 31, 2020**

(54) **METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT**

7,143,300 B2 11/2006 Potter et al.
 7,647,516 B2 1/2010 Ranganathan et al.
 (Continued)

FOREIGN PATENT DOCUMENTS

(71) Applicant: **Lancium LLC**, Houston, TX (US)

CN 103163904 A 6/2013
 KR 20090012523 A 2/2009
 WO 2015199629 A1 12/2015

(72) Inventors: **Michael T. McNamara**, Newport Beach, CA (US); **Raymond E. Cline, Jr.**, Houston, TX (US)

OTHER PUBLICATIONS

(73) Assignee: **Lancium LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Bird et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," National Renewable Energy Lab (NREL), Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.

(Continued)

(21) Appl. No.: **16/702,931**

Primary Examiner — Christopher E. Everett

(22) Filed: **Dec. 4, 2019**

(74) *Attorney, Agent, or Firm* — McDonnell Boehnen Hulbert & Berghoff LLP

Related U.S. Application Data

(60) Provisional application No. 62/927,119, filed on Oct. 28, 2019.

(57) **ABSTRACT**

(51) **Int. Cl.**
H02J 3/14 (2006.01)
H02J 3/00 (2006.01)
G06F 1/3203 (2019.01)

Examples relate to adjusting load power consumption based on a power option agreement. A computing system may receive power option data that is based on a power option agreement and specify minimum power thresholds associated with time intervals. The computing system may determine a performance strategy for a load (e.g., set of computing systems) based on a combination of the power option data and one or more monitored conditions. The performance strategy may specify a power consumption target for the load for each time interval such that each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The computing system may provide instructions the set of computing systems to perform one or more computational operations based on the performance strategy.

(52) **U.S. Cl.**
 CPC **H02J 3/14** (2013.01); **G06F 1/3203** (2013.01); **H02J 3/008** (2013.01)

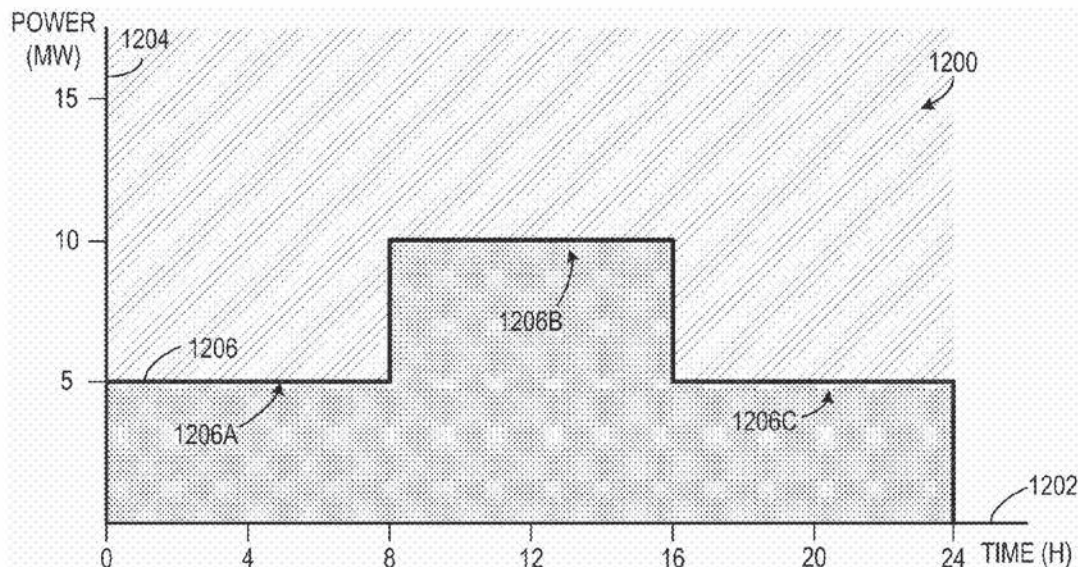
(58) **Field of Classification Search**
 CPC H02J 3/14; H02J 3/008; G06F 1/3203
 See application file for complete search history.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,288,456 B1 9/2001 Cratty
 6,633,823 B2 10/2003 Bartone et al.

20 Claims, 16 Drawing Sheets



US 10,608,433 B1

Page 2

(56)

References Cited

U.S. PATENT DOCUMENTS

7,702,931	B2	4/2010	Goodrum et al.	
7,779,276	B2	8/2010	Bolan et al.	
7,861,102	B1	12/2010	Ranganathan et al.	
7,921,315	B2	4/2011	Langgood et al.	
7,970,561	B2	6/2011	Pfeiffer	
8,001,403	B2	8/2011	Hamilton et al.	
8,006,108	B2	8/2011	Brey et al.	
8,214,843	B2	7/2012	Boss et al.	
8,374,928	B2	2/2013	Gopisetty et al.	
8,447,993	B2	5/2013	Greene et al.	
8,571,820	B2	10/2013	Pfeiffer	
8,627,123	B2	1/2014	Jain et al.	
8,700,929	B1*	4/2014	Weber	G06F 30/13 713/310
8,789,061	B2	7/2014	Pavel et al.	
8,799,690	B2	8/2014	Dawson et al.	
9,003,211	B2	4/2015	Pfeiffer	
9,003,216	B2	4/2015	Sankar et al.	
9,026,814	B2	5/2015	Aasheim et al.	
9,207,993	B2	12/2015	Jain	
9,218,035	B2	12/2015	Li et al.	
9,552,234	B2	1/2017	Boldyrev et al.	
10,367,353	B1	7/2019	McNamara et al.	
10,367,535	B2	7/2019	Corse et al.	
10,444,818	B1	10/2019	McNamara et al.	
10,452,127	B1	10/2019	McNamara et al.	
10,497,072	B2	12/2019	Hooshmand et al.	
2002/0072868	A1	6/2002	Bartone et al.	
2003/0074464	A1	4/2003	Bohrer et al.	
2005/0203761	A1*	9/2005	Barr	G06F 1/26 713/320
2006/0161765	A1	7/2006	Cromer et al.	
2008/0030078	A1	2/2008	Whitted et al.	
2008/0094797	A1	4/2008	Coglitore et al.	
2009/0055665	A1	2/2009	Maglione et al.	
2009/0070611	A1*	3/2009	Bower, III	G06F 1/3203 713/322
2009/0089595	A1*	4/2009	Brey	G06F 1/3203 713/300
2010/0211810	A1	8/2010	Zacho	
2010/0328849	A1	12/2010	Ewing et al.	
2011/0072289	A1	3/2011	Kato	
2012/0000121	A1	1/2012	Swann	
2012/0072745	A1	3/2012	Ahluwalia et al.	
2012/0300524	A1	11/2012	Fornage et al.	
2012/0324259	A1	12/2012	Aasheim et al.	
2013/0006401	A1	1/2013	Shan	
2013/0063991	A1	3/2013	Xiao et al.	
2013/0086404	A1*	4/2013	Sankar	G06F 1/305 713/324
2013/0187464	A1	7/2013	Smith et al.	
2013/0227139	A1	8/2013	Suffling	
2013/0306276	A1	11/2013	Duchesneau	
2014/0137468	A1	5/2014	Ching	
2014/0379156	A1	12/2014	Kamel et al.	
2015/0121113	A1	4/2015	Ramamurthy et al.	
2015/0155712	A1	6/2015	Mondal	
2015/0229227	A1	8/2015	Aeloiza et al.	
2015/0277410	A1	10/2015	Gupta et al.	
2015/0278968	A1	10/2015	Steven et al.	
2016/0170469	A1	6/2016	Schgal et al.	
2016/0172900	A1	6/2016	Welch, Jr.	
2016/0187906	A1*	6/2016	Bodas	G05B 15/02 700/287

2016/0198656	A1	7/2016	McNamara et al.
2016/0212954	A1	7/2016	Argento
2016/0324077	A1	11/2016	Frantzen et al.
2017/0023969	A1	1/2017	Shows et al.
2017/0104336	A1	4/2017	Elbsat et al.
2017/0261949	A1	9/2017	Hoffmann et al.
2018/0144414	A1	5/2018	Lee et al.
2018/0202825	A1	7/2018	You et al.
2018/0240112	A1	8/2018	Castinado et al.
2018/0366978	A1	12/2018	Matan et al.
2018/0367320	A1	12/2018	Montalvo
2019/0052094	A1	2/2019	Pmsvsv et al.
2019/0168630	A1	6/2019	Mrlik et al.
2019/0258307	A1	8/2019	Shaikh et al.
2019/0280521	A1	9/2019	Lundstrom et al.
2019/0318327	A1	10/2019	Sowell et al.
2019/0324820	A1	10/2019	Krishnan et al.

OTHER PUBLICATIONS

EPEX Spot, "How They Occur, What They Mean," https://www.epexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.

Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.

Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017955, dated Apr. 30, 2018, 22 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017950, dated May 31, 2018, 15 pages.

Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.

Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.

Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.

Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.

Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.

Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.

Rahimi, Farrokh, "Using a Transactive Energy Framework," IEEE Electrification Magazine, Dec. 2016, pp. 23-29.

Soluna, "Powering the Block Chain," Aug. 2018, version 1.1, 29 pages.

Wilson, Joseph Nathanael, "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016, 154 pages.

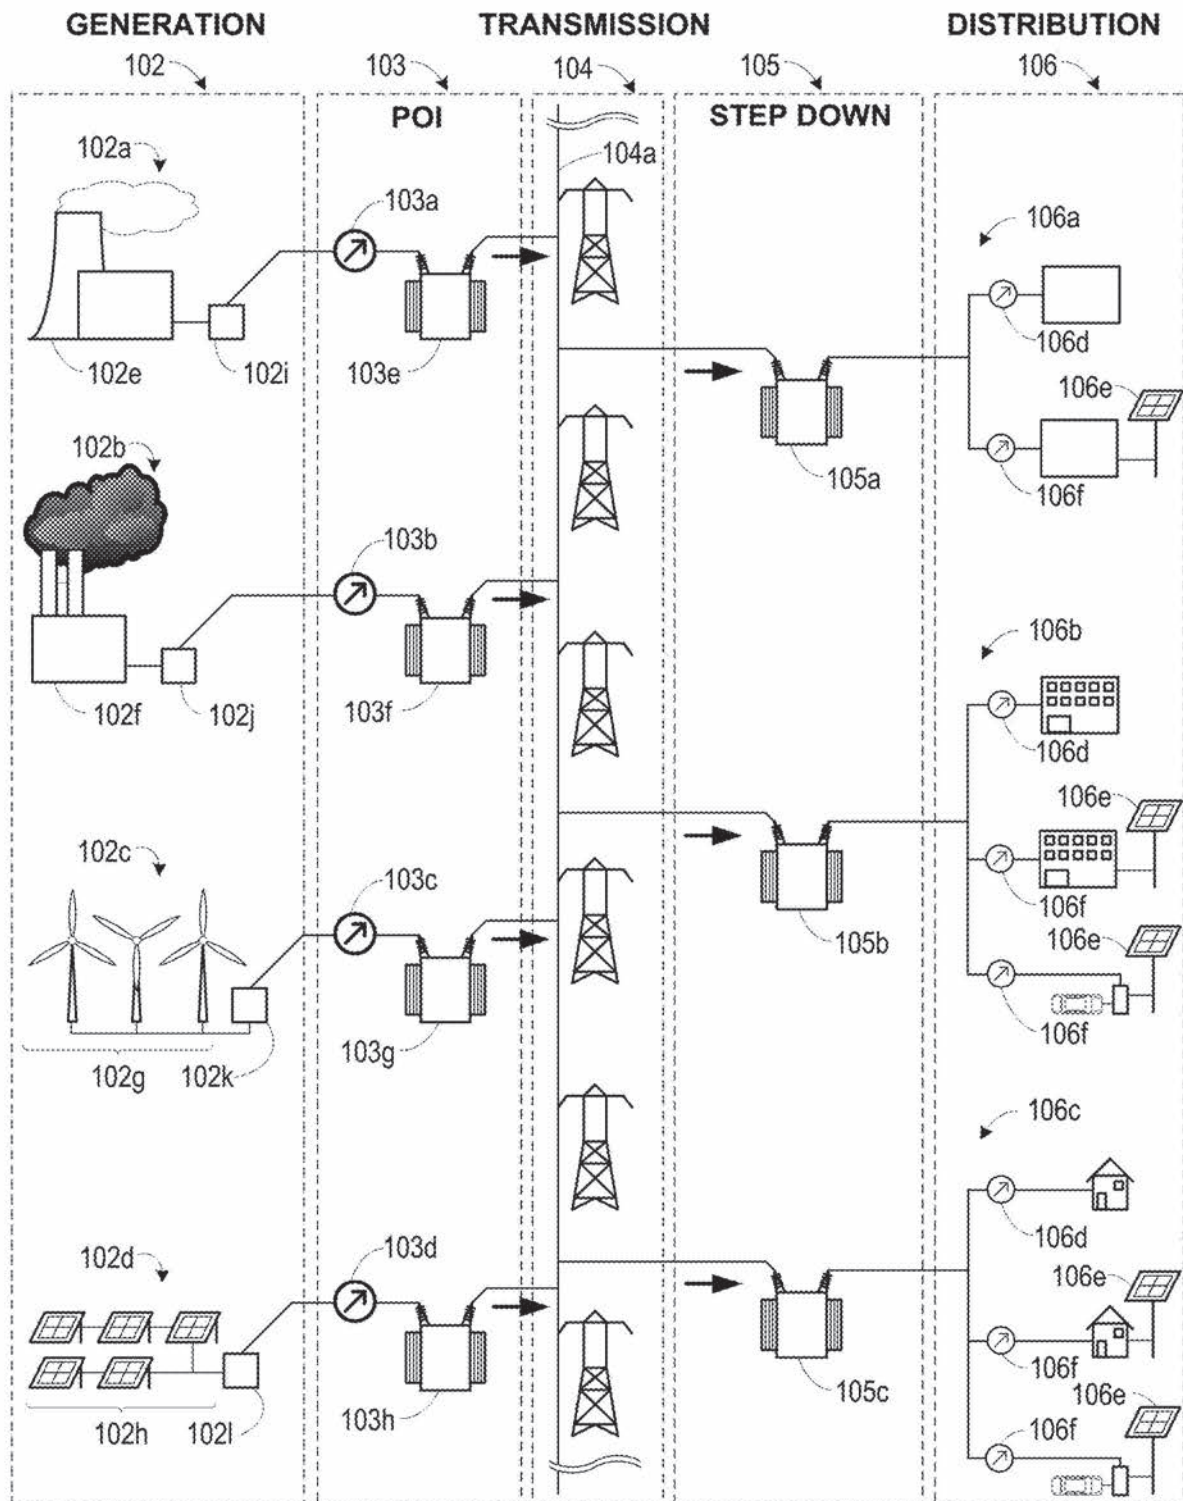
* cited by examiner

U.S. Patent

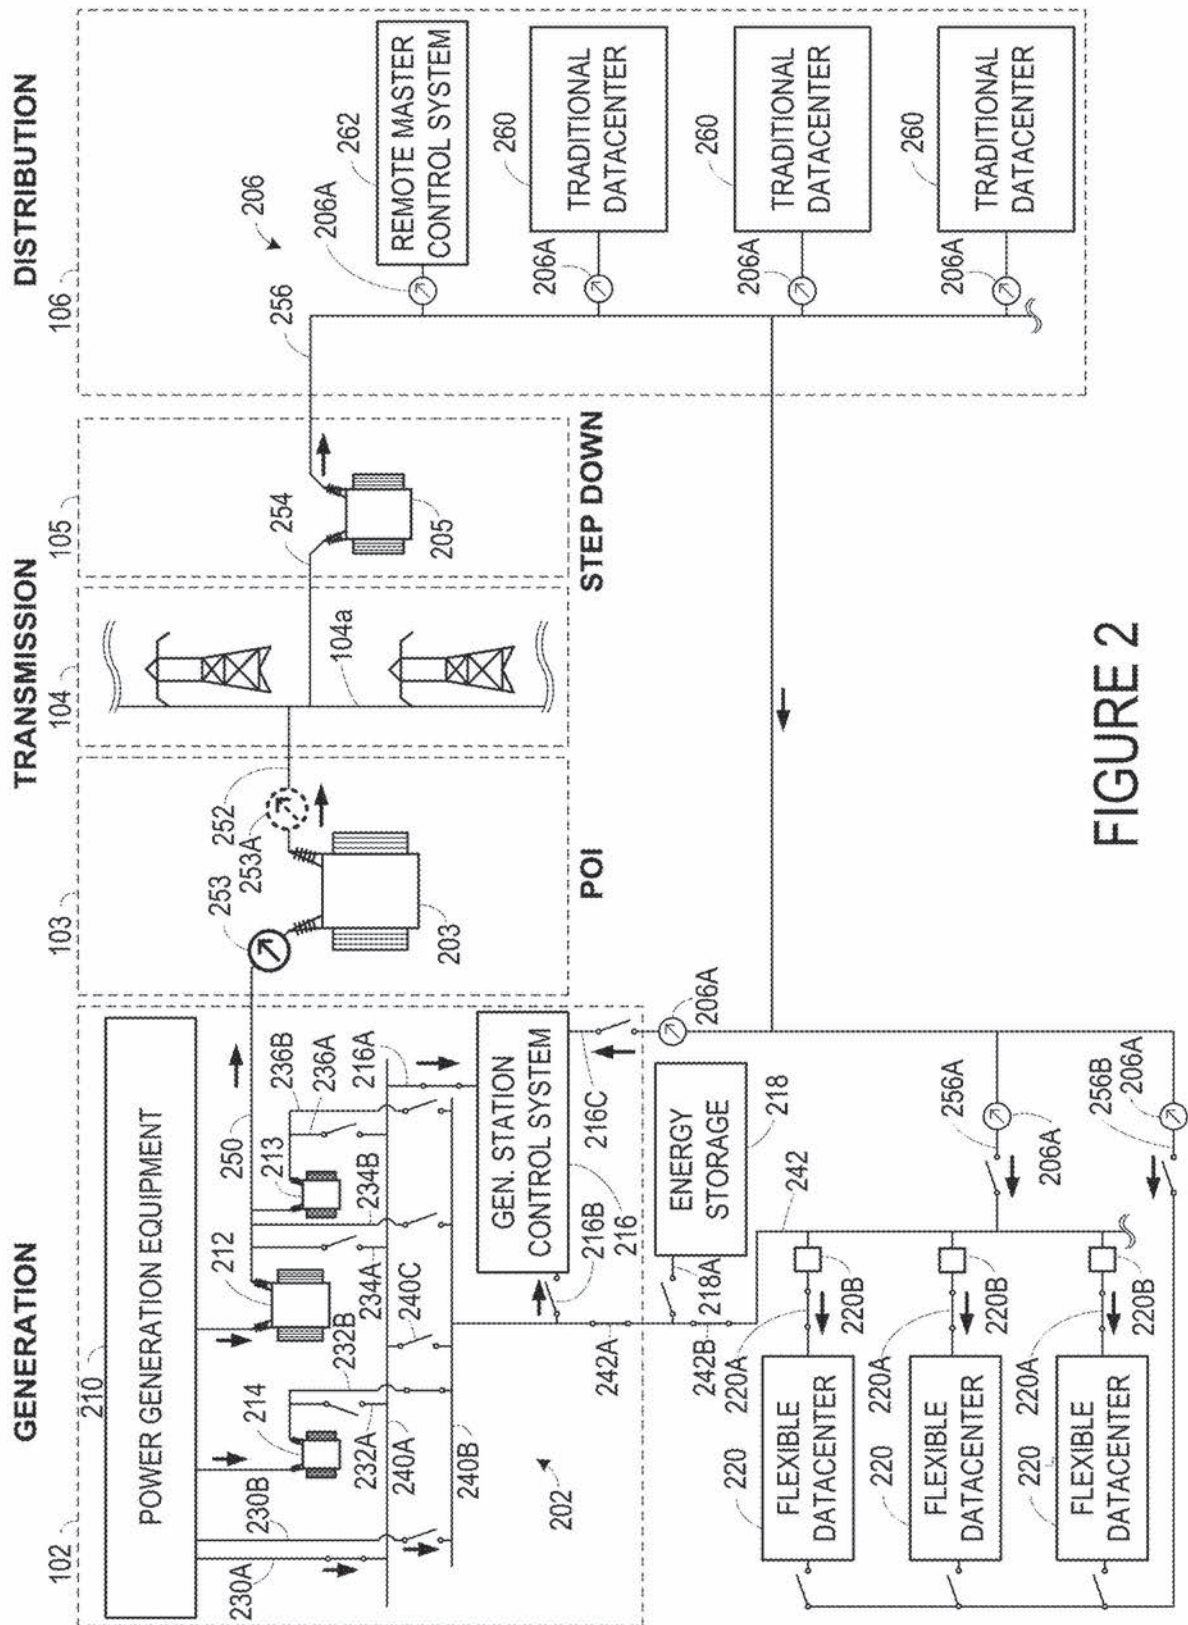
Mar. 31, 2020

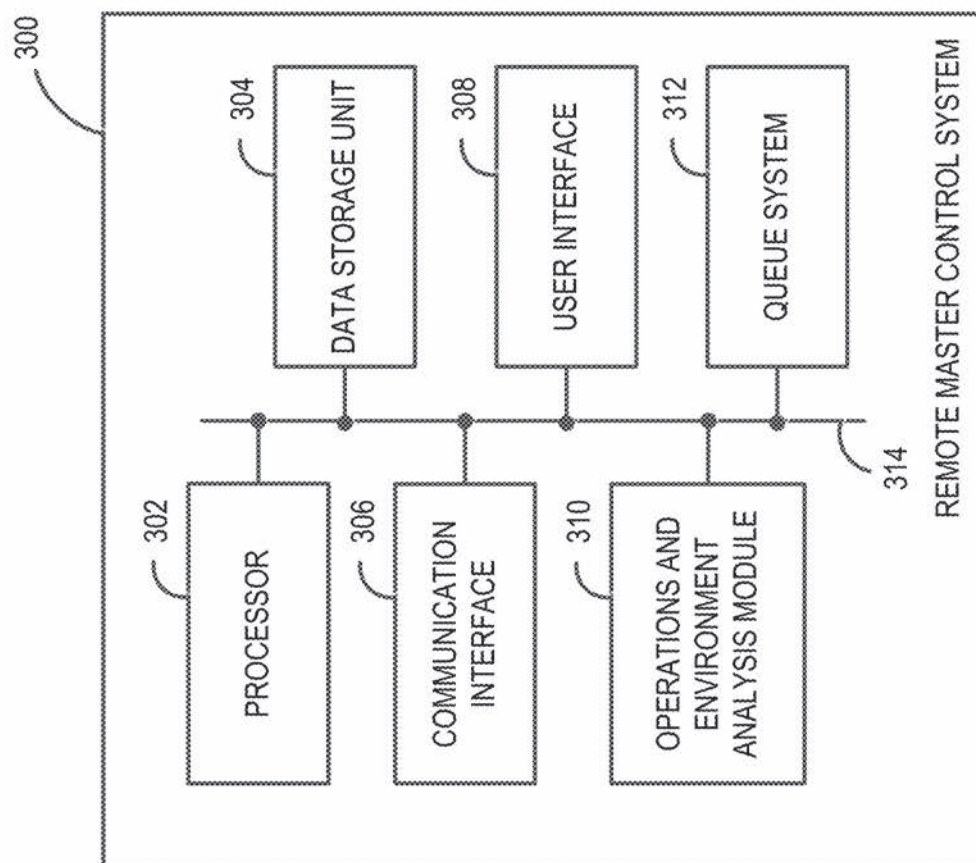
Sheet 1 of 16

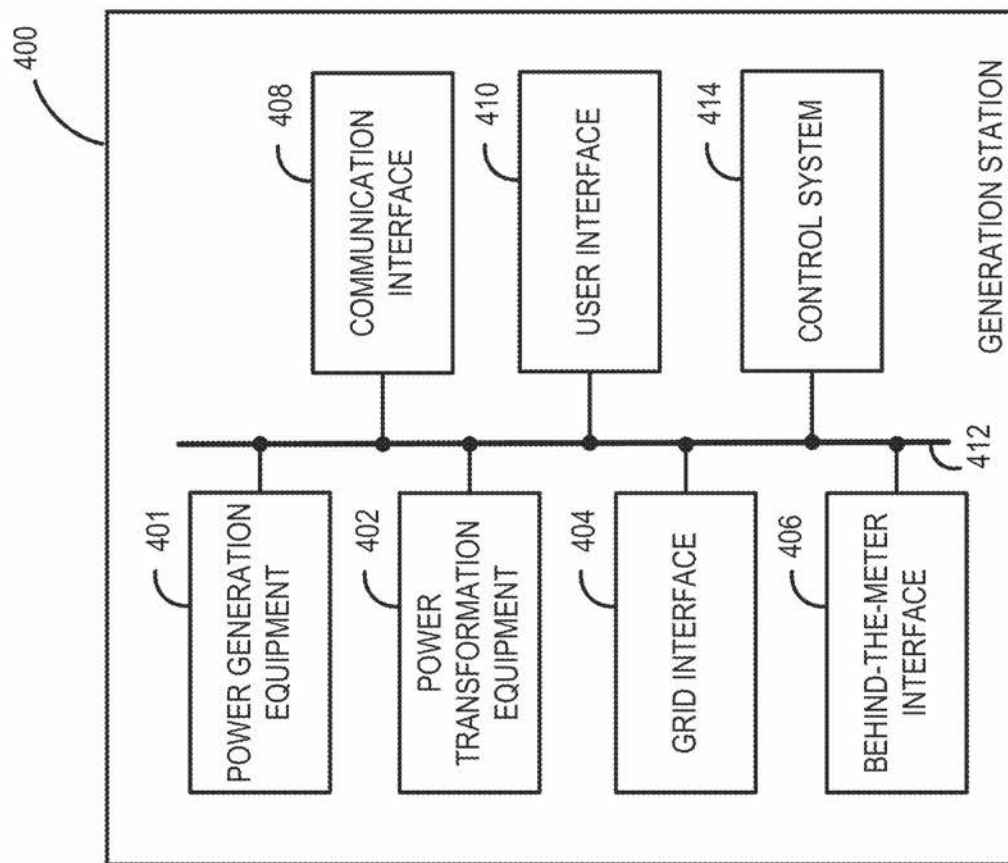
US 10,608,433 B1



PRIOR ART
FIGURE 1







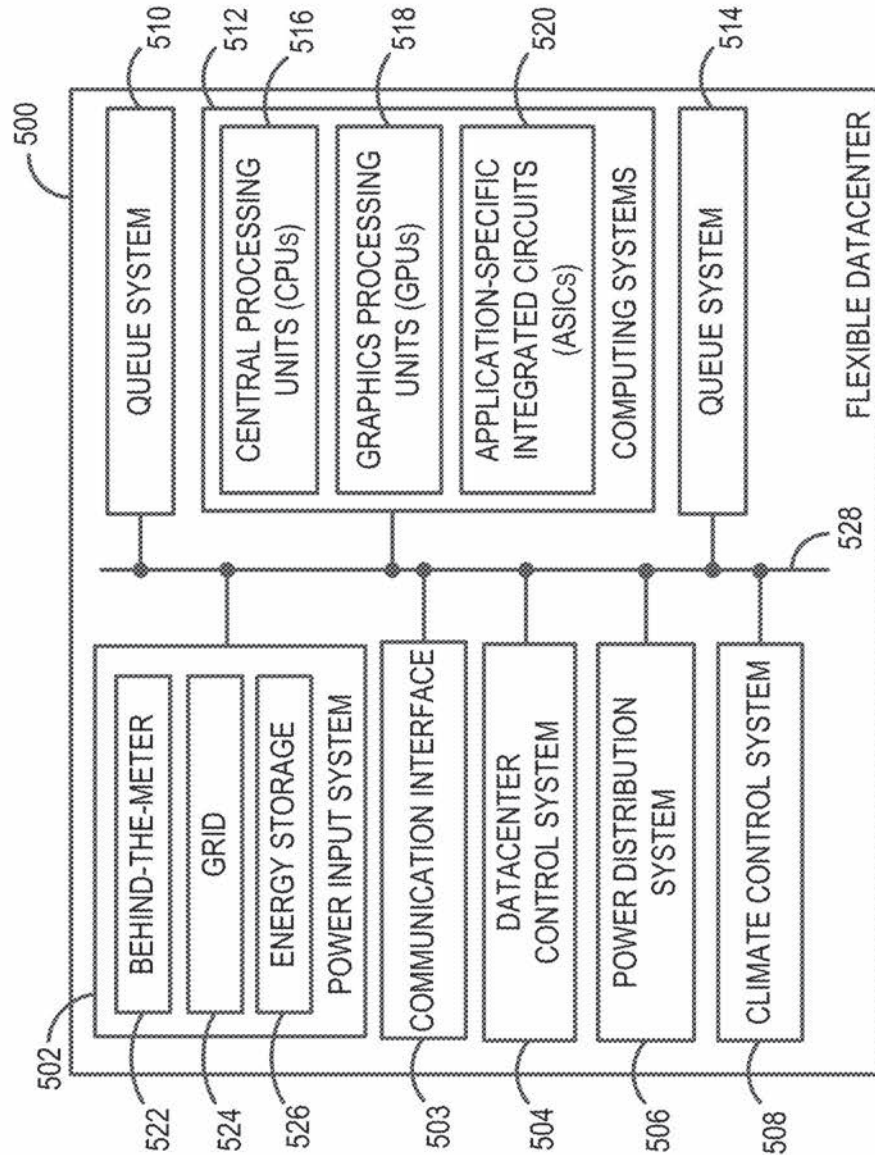
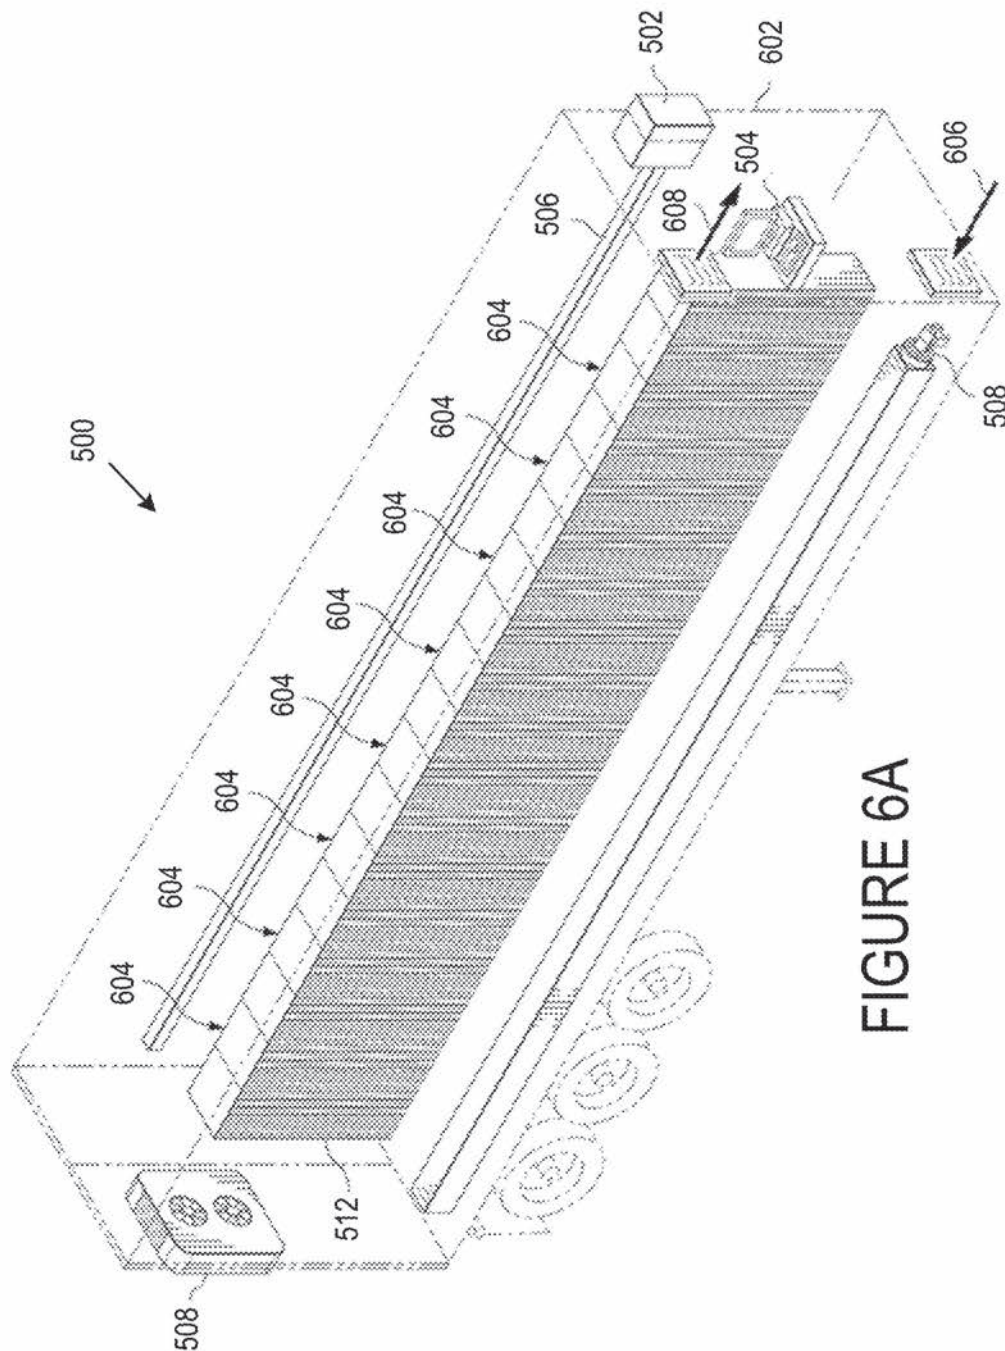


FIGURE 5



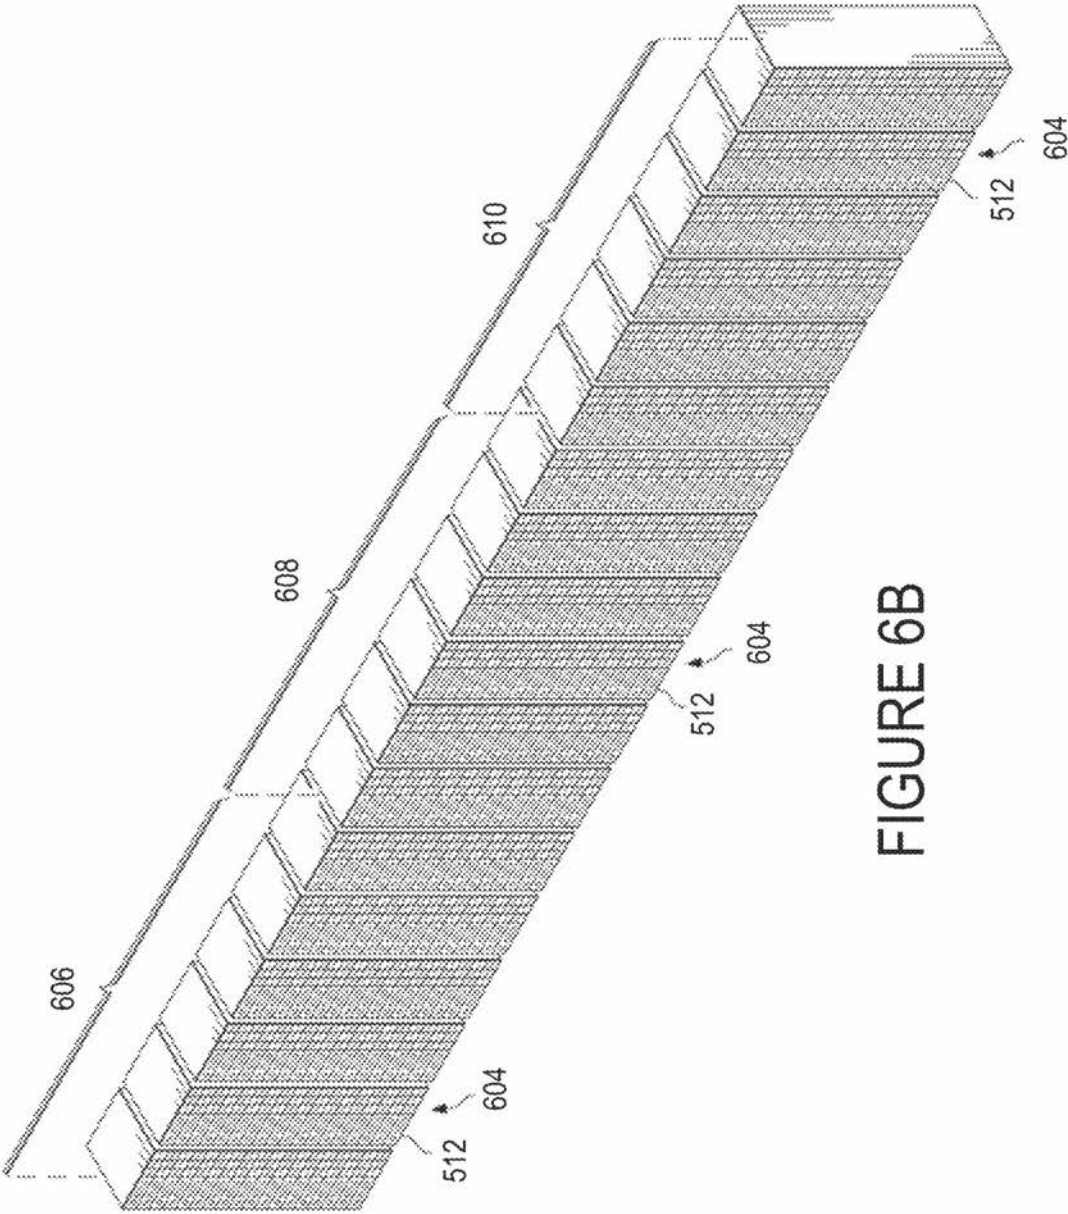


FIGURE 6B

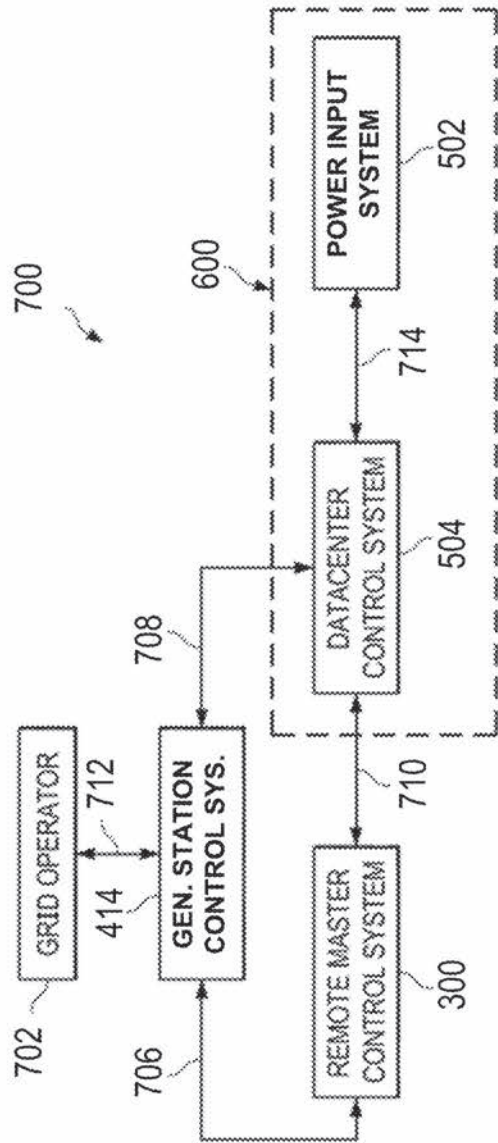
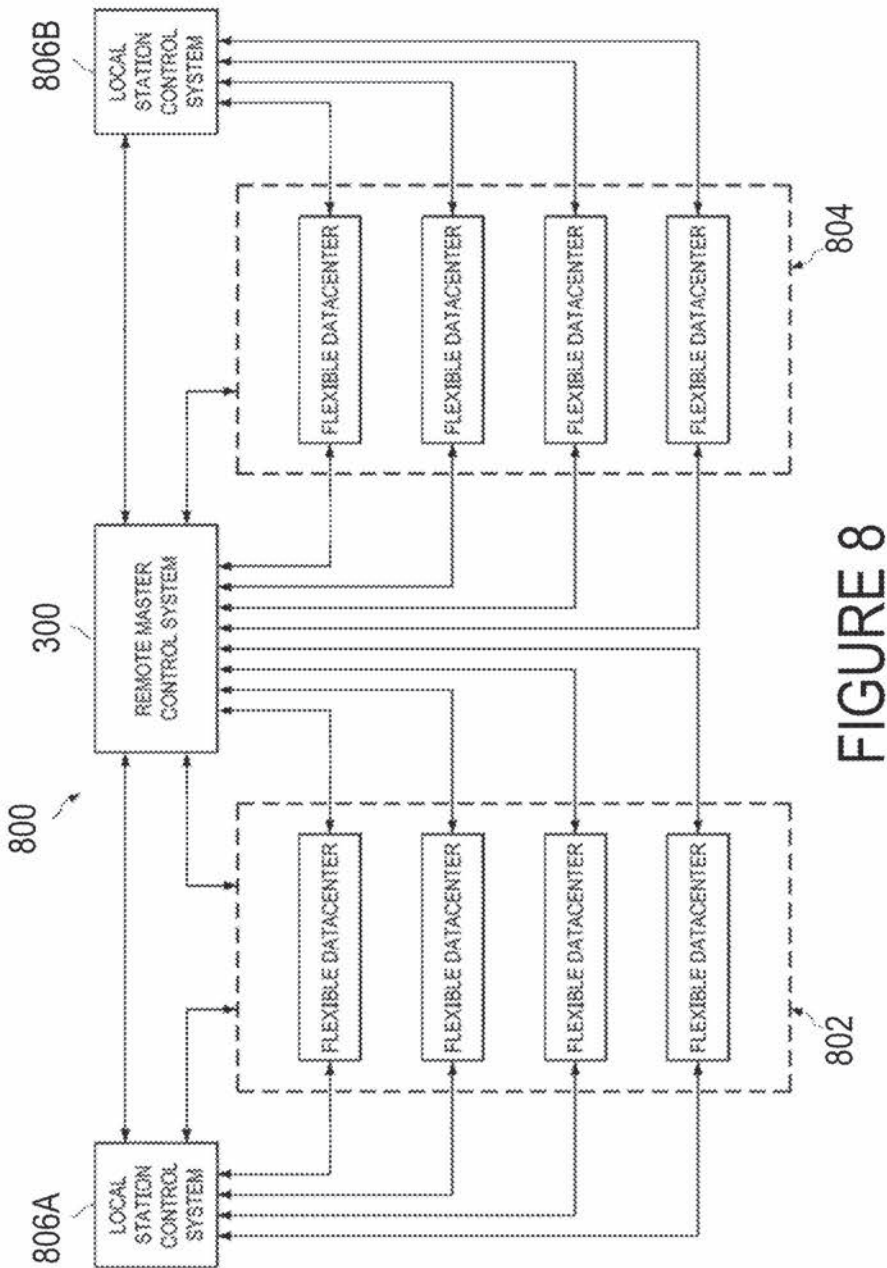
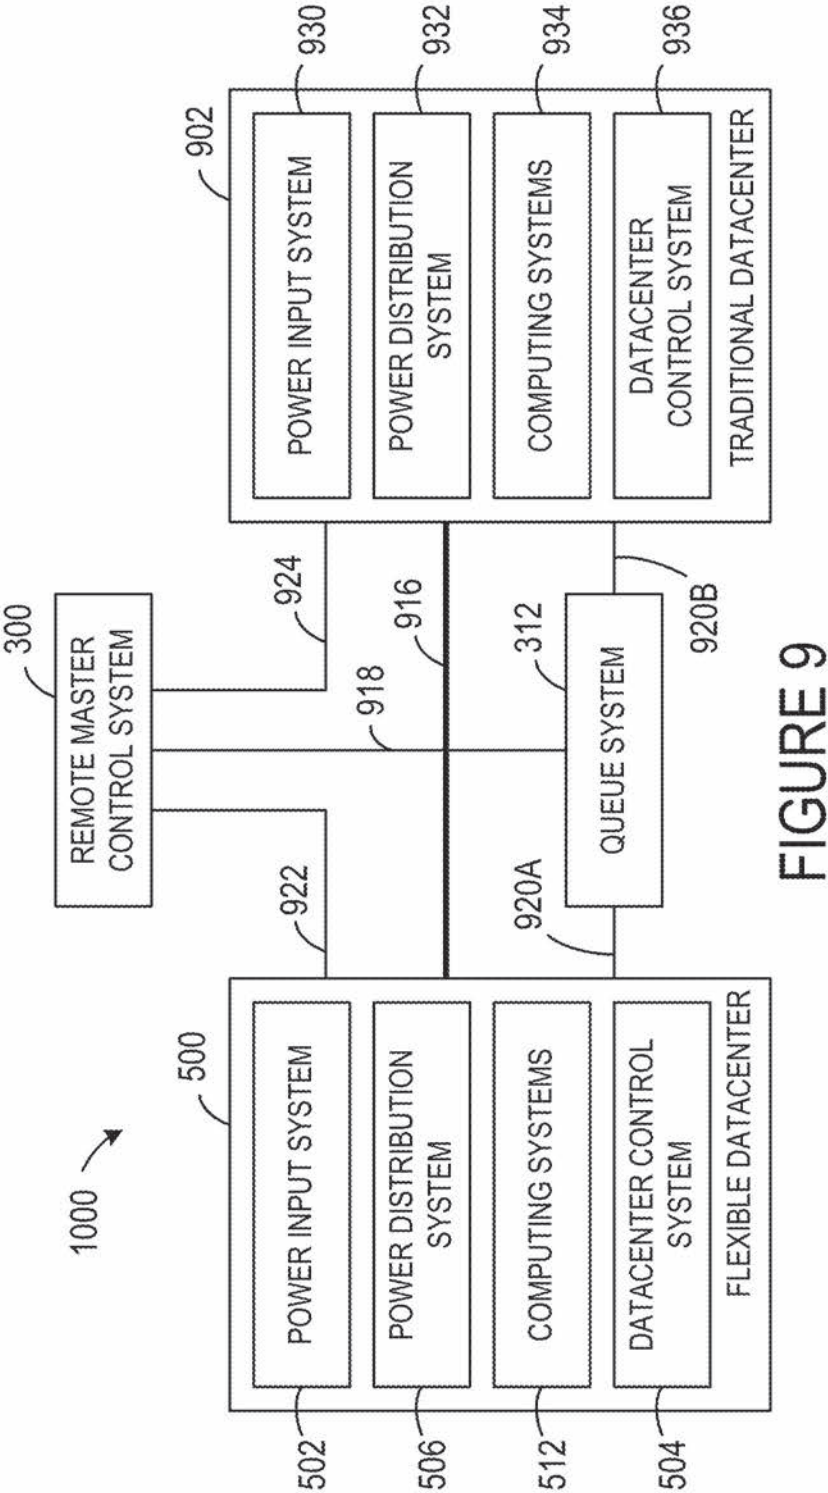


FIGURE 7





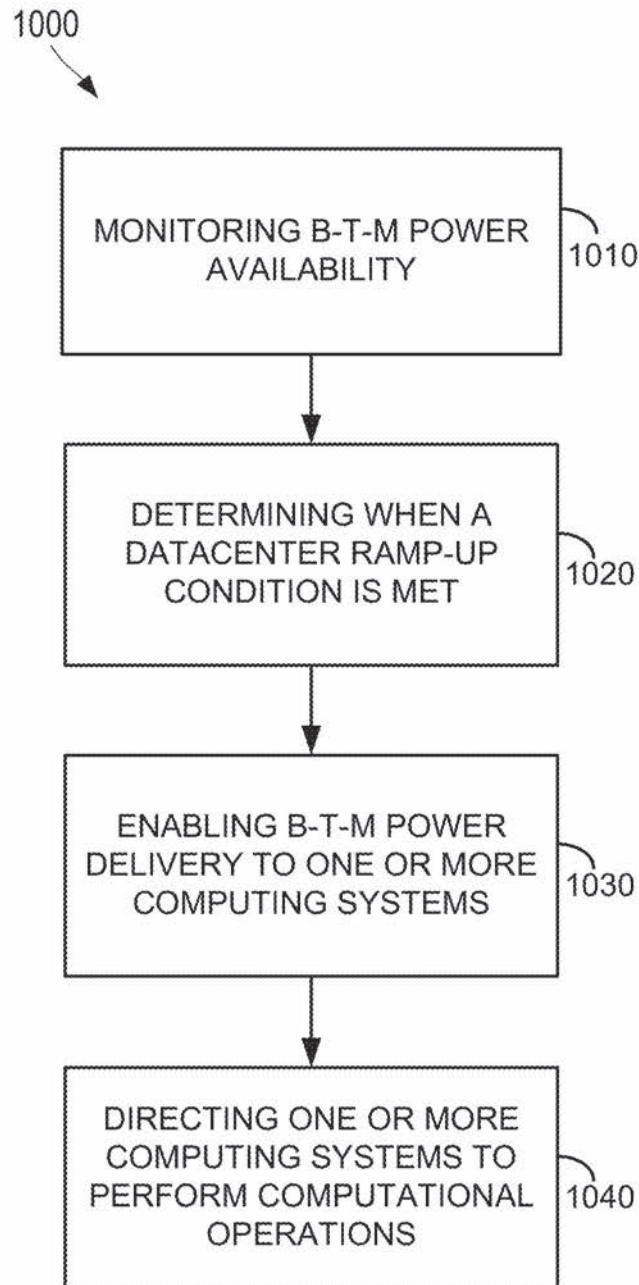


FIGURE 10A

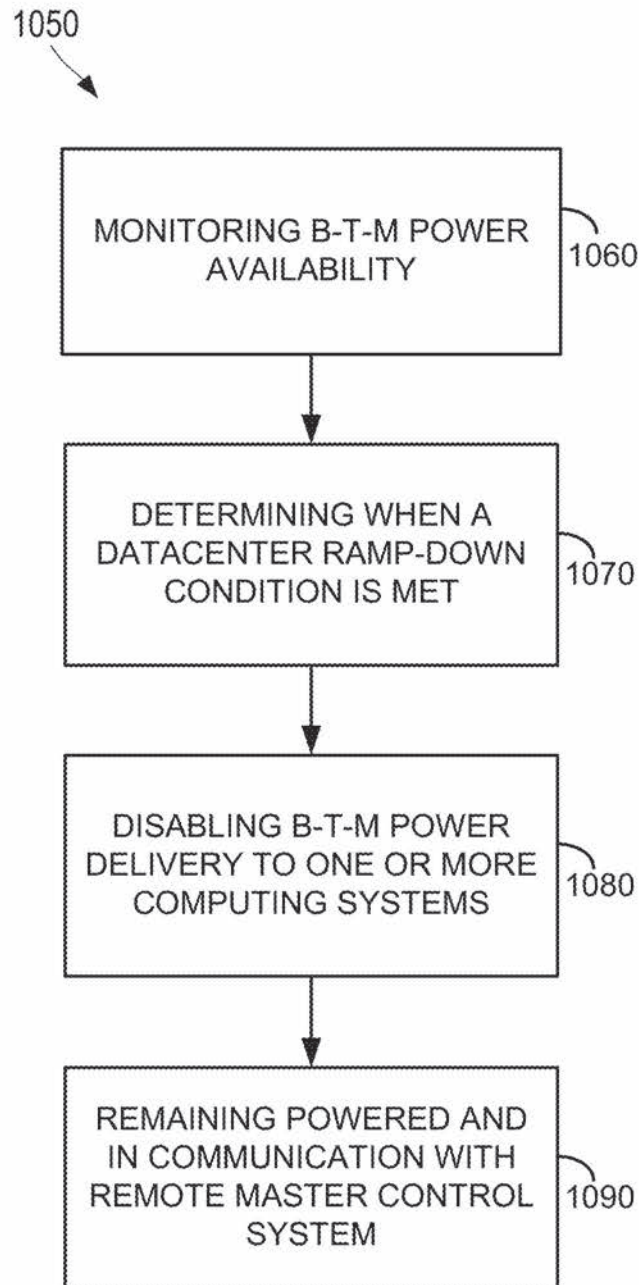


FIGURE 10B

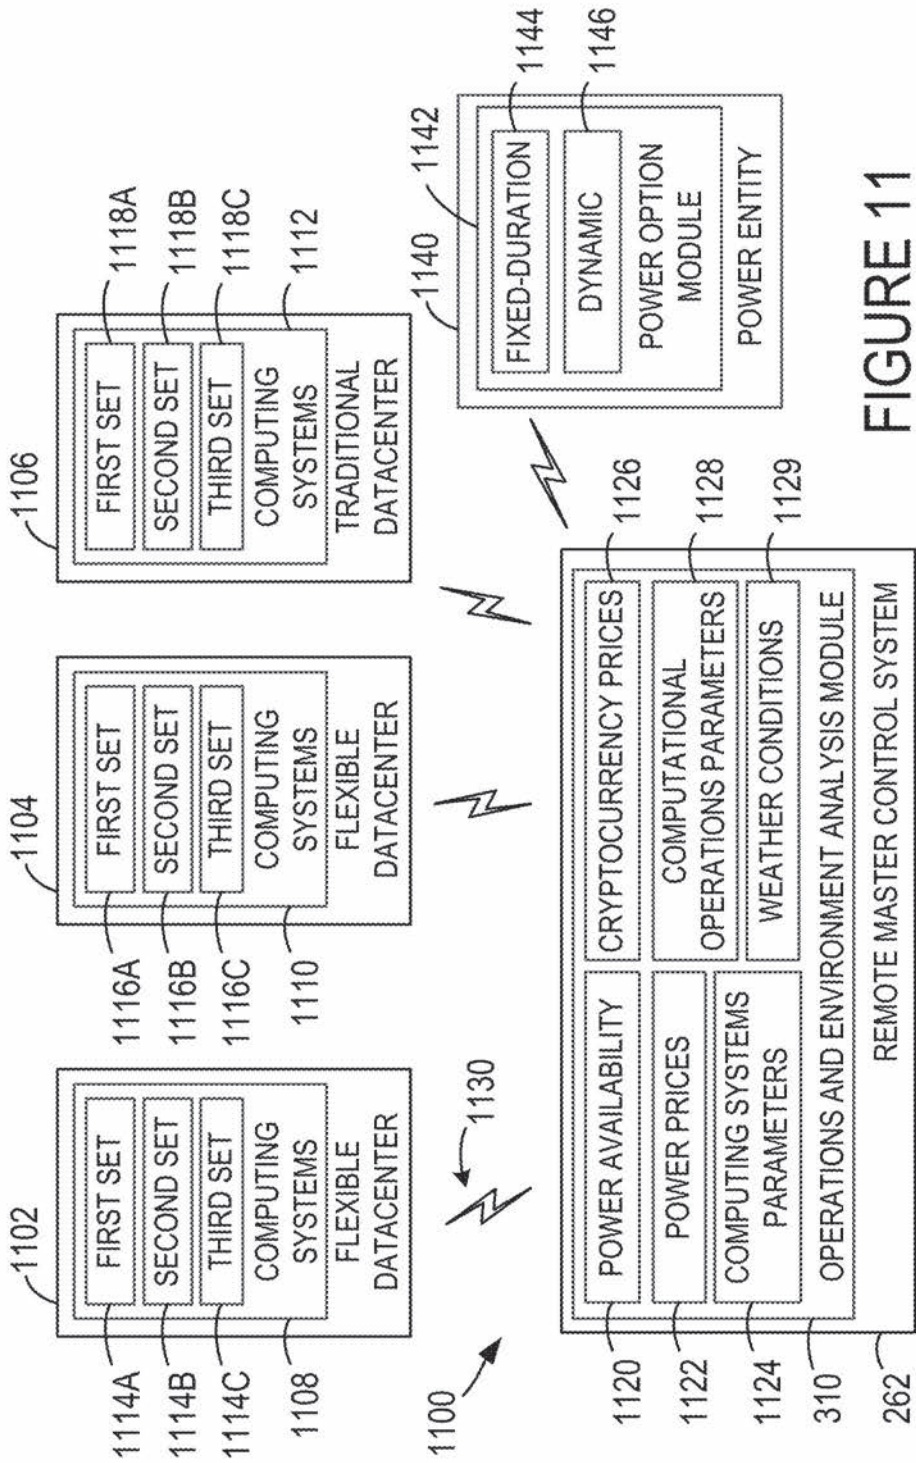


FIGURE 11

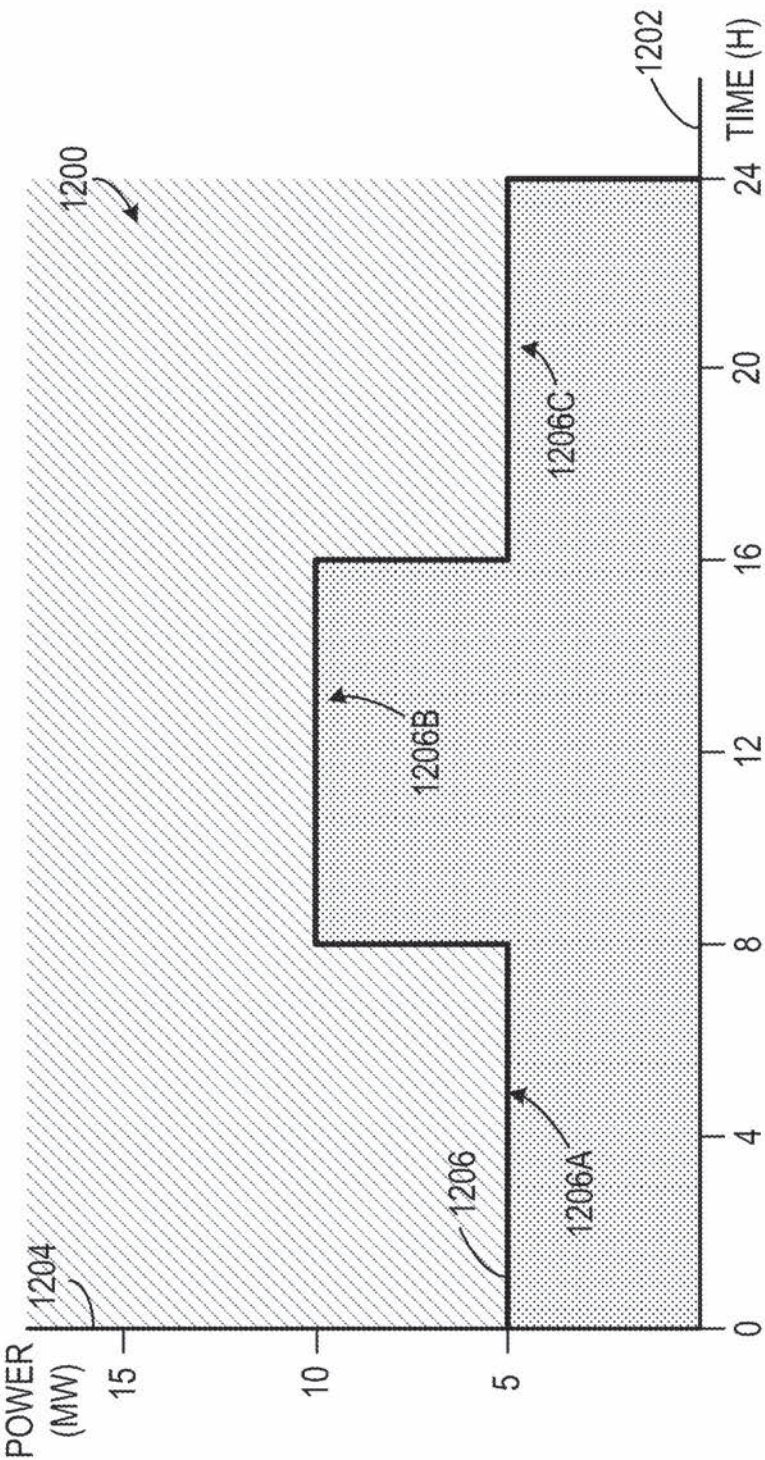


FIGURE 12

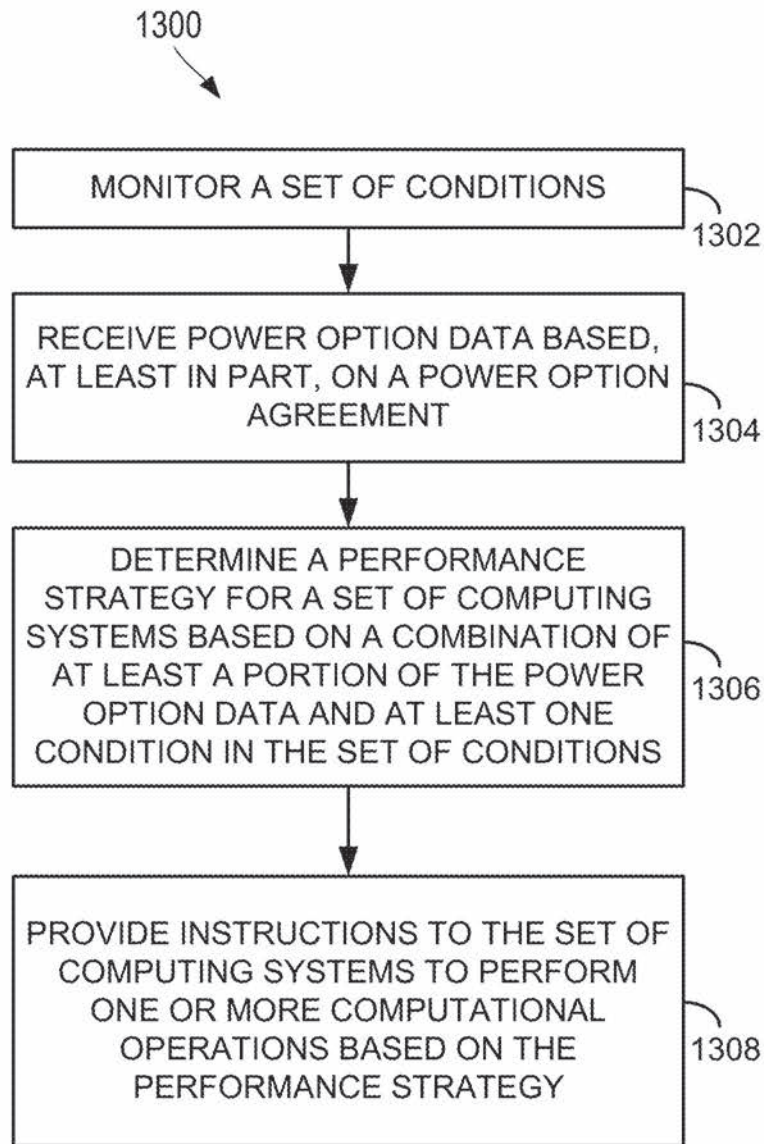


FIGURE 13

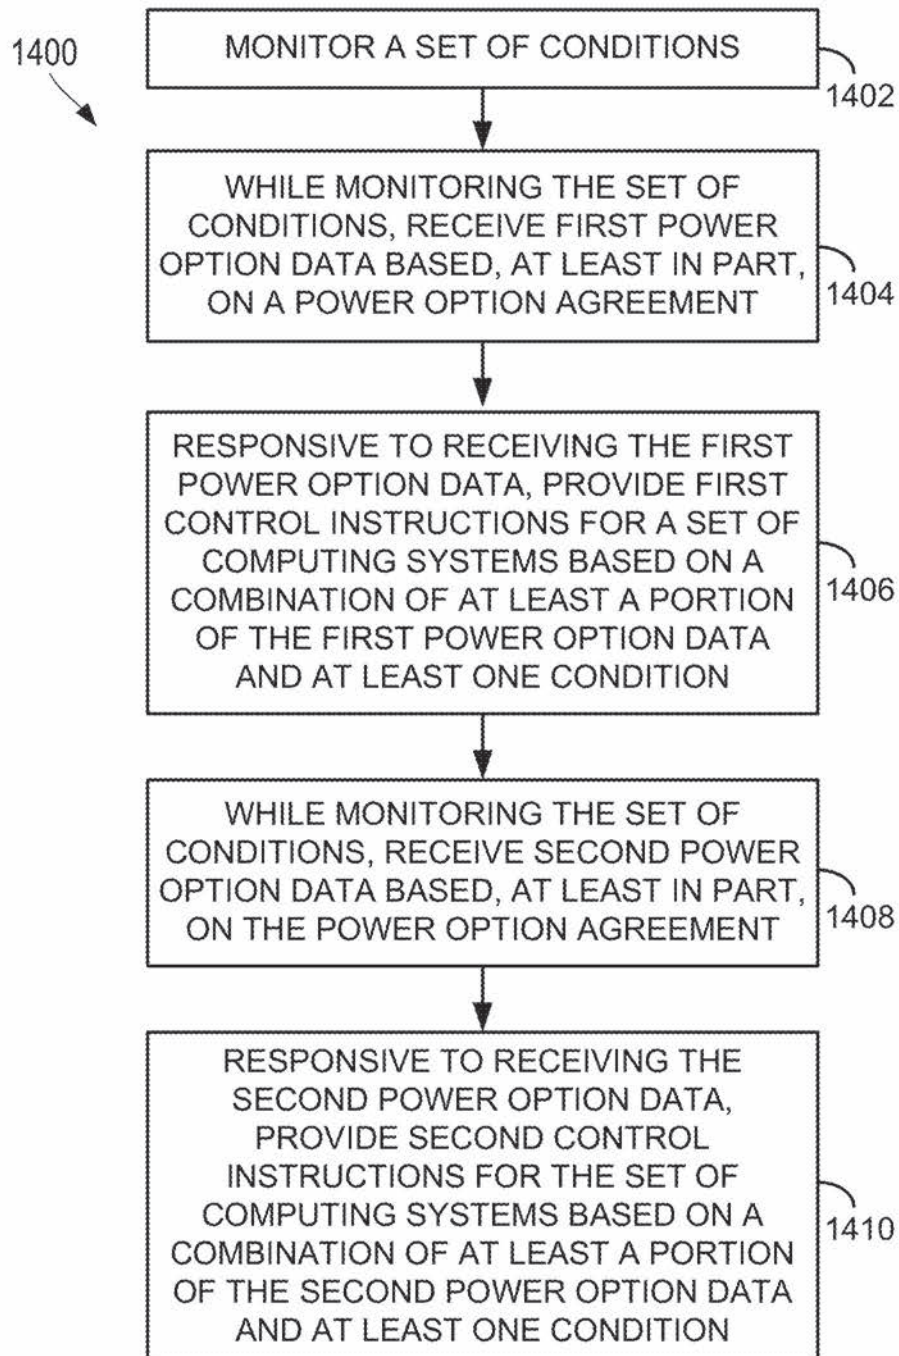


FIGURE 14

US 10,608,433 B1

1

METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Patent Application No. 62/927,119, filed Oct. 28, 2019, the entire contents of which are herein incorporated by reference.

FIELD

This specification relates to power consumption adjustments when using grid power and/or intermittent behind-the-meter power.

BACKGROUND

“Electrical grid” or “grid,” as used herein, refers to a Wide Area Synchronous Grid (also known as an Interconnection), and is a regional scale or greater electric power grid that operates at a synchronized frequency and is electrically tied together during normal system conditions. An electrical grid delivers electricity from generation stations to consumers. An electrical grid includes: (i) generation stations that produce electrical power at large scales for delivery through the grid, (ii) high voltage transmission lines that carry that power from the generation stations to demand centers, and (iii) distribution networks carry that power to individual customers.

FIG. 1 illustrates a typical electrical grid, such as a North American Interconnection or the synchronous grid of Continental Europe (formerly known as the UCTE grid). The electrical grid of FIG. 1 can be described with respect to the various segments that make up the grid.

A generation segment **102** includes one or more generation stations that produce utility-scale electricity (typically >50 MW), such as a nuclear plant **102a**, a coal plant **102b**, a wind power station (i.e., wind farm) **102c**, and/or a photovoltaic power station (i.e., a solar farm) **102d**. Generation stations are differentiated from building-mounted and other decentralized or local wind or solar power applications because they supply power at the utility level and scale (>50 MW), rather than to a local user or users. The primary purpose of generation stations is to produce power for distribution through the grid, and in exchange for payment for the supplied electricity. Each of the generation stations **102a-d** includes power generation equipment **102e-h**, respectively, typically capable of supply utility-scale power (>50 MW). For example, the power generation equipment **102g** at wind power station **102c** includes wind turbines, and the power generation equipment **102h** at photovoltaic power station **102d** includes photovoltaic panels.

Each of the generation stations **102a-d** may further include station electrical equipment **102i-1** respectively. Station electrical equipment **102i-1** are each illustrated in FIG. 1 as distinct elements for simplified illustrative purposes only and may, alternatively or additionally, be distributed throughout the power generation equipment, **102e-h**, respectively. For example, at wind power station **102c**, each wind turbine may include transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each wind turbine may be collected by distribution lines along strings of wind turbines and move through

2

collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the wind power station **102c**. Similarly, at photovoltaic power station **102d**, individual photovoltaic panels and/or arrays of photovoltaic panels may include inverters, transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each photovoltaic panel and/or array may be collected by distribution lines along the photovoltaic panels and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the photovoltaic power station **102d**.

Each generation station **102a-d** may produce AC or DC electrical current which is then typically stepped up to a higher AC voltage before leaving the respective generation station. For example, wind turbines may typically produce AC electrical energy at 600V to 700V, which may then be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102c**. As another example, photovoltaic arrays may produce DC voltage at 600V to 900V, which is then inverted to AC voltage and may be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102d**.

Upon exiting the generation segment **102**, electrical power generated at generation stations **102a-d** passes through a respective Point of Interconnection (“POI”) **103** between a generation station (e.g., **102a-d**) and the rest of the grid. A respective POI **103** represents the point of connection between a generation station’s (e.g., **102a-d**) equipment and a transmission system (e.g., transmission segment **104**) associated with electrical grid. In some cases, at the POI **103**, generated power from generation stations **102a-d** may be stepped up at transformer systems **103e-h** to high voltage scales suitable for long-distance transmission along transmission lines **104a**. Typically, the generated electrical energy leaving the POI **103** will be at 115 kV AC or above, but in some cases it may be as low as, for example, 69 kV for shorter distance transmissions along transmission lines **104a**. Each of transformer systems **103e-h** may be a single transformer or may be multiple transformers operating in parallel or series and may be co-located or located in geographically distinct locations. Each of the transformer systems **103e-h** may include substations and other links between the generation stations **102a-d** and the transmission lines **104a**.

A key aspect of the POI **103** is that this is where generation-side metering occurs. One or more utility-scale generation-side meters **103a-d** (e.g., settlement meters) are located at settlement metering points at the respective POI **103** for each generation station **102a-d**. The utility-scale generation-side meters **103a-d** measure power supplied from generation stations **102a-d** into the transmission segment **104** for eventual distribution throughout the grid.

For electricity consumption, the price consumers pay for power distributed through electric power grids is typically composed of, among other costs, Generation, Administration, and Transmission & Distribution (“T&D”) costs. T&D costs represent a significant portion of the overall price paid by consumers for electricity. These costs include capital costs (land, equipment, substations, wire, etc.), costs associated with electrical transmission losses, and operation and maintenance costs.

US 10,608,433 B1

3

For utility-scale electricity supply, operators of generation stations (e.g., 102a-d) are paid a variable market price for the amount of power the operator generates and provides to the grid, which is typically determined via a power purchase agreement (PPA) between the generation station operator and a grid operator. The amount of power the generation station operator generates and provides to the grid is measured by utility-scale generation-side meters (e.g., 103a-d) at settlement metering points. As illustrated in FIG. 1, the utility-scale generation-side meters 103a-d are shown on a low side of the transformer systems 103e-h, but they may alternatively be located within the transformer systems 103e-h or on the high side of the transformer systems 103e-h. A key aspect of a utility-scale generation-side meter is that it is able to meter the power supplied from a specific generation station into the grid. As a result, the grid operator can use that information to calculate and process payments for power supplied from the generation station to the grid. That price paid for the power supplied from the generation station is then subject to T&D costs, as well as other costs, in order to determine the price paid by consumers.

After passing through the utility-scale generation-side meters in the POI 103, the power originally generated at the generation stations 102a-d is transmitted onto and along the transmission lines 104a in the transmission segment 104. Typically, the electrical energy is transmitted as AC at 115 kV+ or above, though it may be as low as 69 kV for short transmission distances. In some cases, the transmission segment 104 may include further power conversions to aid in efficiency or stability. For example, transmission segment 104 may include high-voltage DC ("HVDC") portions (along with conversion equipment) to aid in frequency synchronization across portions of the transmission segment 104. As another example, transmission segment 104 may include transformers to step AC voltage up and then back down to aid in long distance transmission (e.g., 230 kV, 500 kV, 765 kV, etc.).

Power generated at the generation stations 104a-d is ultimately destined for use by consumers connected to the grid. Once the energy has been transmitted along the transmission segment 104, the voltage will be stepped down by transformer systems 105a-c in the step down segment 105 so that it can move into the distribution segment 106.

In the distribution segment 106, distribution networks 106a-c take power that has been stepped down from the transmission lines 104a and distribute it to local customers, such as local sub-grids (illustrated at 106a), industrial customers, including large EV charging networks (illustrated at 106b), and/or residential and retail customers, including individual EV charging stations (illustrated at 106c). Customer meters 106d, 106f measure the power used by each of the grid-connected customers in distribution networks 106a-c. Customer meters 106d are typically load meters that are unidirectional and measure power use. Some of the local customers in the distribution networks 106a-d may have local wind or solar power systems 106e owned by the customer. As discussed above, these local customer power systems 106e are decentralized and supply power directly to the customer(s). Customers with decentralized wind or solar power systems 106e may have customer meters 106f that are bidirectional or net-metering meters that can track when the local customer power systems 106e produce power in excess of the customer's use, thereby allowing the utility to provide a credit to the customer's monthly electricity bill. Customer meters 106d, 106f differ from utility-scale generation-side meters (e.g., settlement meters) in at least the following characteristics: design (electro-mechanical or electronic vs

4

current transformer), scale (typically less than 1600 amps vs. typically greater than 50 MW; typically less than 600V vs. typically greater than 14 kV), primary function (use vs. supply metering), economic purpose (credit against use vs. payment for power), and location (in a distribution network at point of use vs. at a settlement metering point at a Point of Interconnection between a generation station and a transmission line).

To maintain stability of the grid, the grid operator strives to maintain a balance between the amount of power entering the grid from generation stations (e.g., 102a-d) and the amount of grid power used by loads (e.g., customers in the distribution segment 106). In order to maintain grid stability and manage congestion, grid operators may take steps to reduce the supply of power arriving from generation stations (e.g., 102a-d) when necessary (e.g., curtailment). Particularly, grid operators may decrease the market price paid for generated power to dis-incentivize generation stations (e.g., 102a-d) from generating and supplying power to the grid. In some cases, the market price may even go negative such that generation station operators must pay for power they allow into the grid. In addition, some situations may arise where grid operators explicitly direct a generation station (e.g., 102a-d) to reduce or stop the amount of power the station is supplying to the grid.

Power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewable energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

According to the National Renewable Energy Laboratory's Technical Report TP-6A20-60983 (March 2014):

[C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission congestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

"Curtailment" today does not necessarily mean what it did in the early 2000s. Two separate changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them. Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being

US 10,608,433 B1

5

replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator’s capital costs are included in a utility’s rate base.

Curtailment may result in available energy being wasted because solar and wind operators have zero variable cost (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

SUMMARY

In an example, a system includes a set of computing systems. The set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and, while monitoring the set of conditions, receive first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The control system is further configured to provide first control instructions for the set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions responsive to receiving the first power option data. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The control system is also configured to, while monitoring the set of conditions, receive second power option data based, at least in part, on the power option

6

agreement. The second power option data specify a second minimum power threshold associated with a second time interval. Responsive to receiving the second power option data, the control system is configured to provide second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and wherein the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In another example, a method involves monitoring, at a computing system, a set of conditions, and while monitoring the set of conditions, receiving first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The method further involves, responsive to receiving the first power option data, providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The method further involves, while monitoring the set of conditions, receiving second power option data based, at least in part, on the power option agreement. The second power option data specify a second minimum power threshold associated with a second time interval. The method also involves, responsive to receiving the second power option data, providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In yet another example, a system is provided. The system includes a set of computing systems, where the set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and receive power option data based, at least in part, on a power option agreement. The power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, where each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals. The control system is further configured to, responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, where each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The control system is also configured to provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.

In a further example, non-transitory computer-readable medium is described that is configured to store instructions,

US 10,608,433 B1

7

that when executed by a computing system, causes the computing system to perform operations consistent with the method steps described above.

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 shows a typical electrical grid.

FIG. 2 shows a behind-the-meter arrangement with optional grid power, including one or more flexible datacenters, according to one or more example embodiments.

FIG. 3 shows a block diagram of a remote master control system, according to one or more example embodiments.

FIG. 4 a block diagram of a generation station, according to one or more example embodiments.

FIG. 5 shows a block diagram of a flexible datacenter, according to one or more example embodiments.

FIG. 6A shows a structural arrangement of a flexible datacenter, according to one or more example embodiments.

FIG. 6B shows a set of computing systems arranged in a straight configuration, according to one or more example embodiments.

FIG. 7 shows a control distribution system for a flexible datacenter, according to one or more example embodiments.

FIG. 8 shows a control distribution system for a fleet of flexible datacenters, according to one or more example embodiments.

FIG. 9 shows a queue distribution system for a traditional datacenter and a flexible datacenter, according to one or more example embodiments.

FIG. 10A shows a method of dynamic power consumption at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 10B shows a method of dynamic power delivery at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 11 shows a block diagram of a system for implementing power consumption adjustments based on a power option agreement, according to one or more embodiments.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments.

FIG. 13 shows a method for implementing power consumption adjustments based on a fixed-duration power option agreement, according to one or more embodiments.

FIG. 14 shows a method for implementing power consumption adjustments based on a dynamic power option agreement, according to one or more embodiments.

DETAILED DESCRIPTION

Disclosed examples will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all of the disclosed examples are shown. Different examples may be described and should not be construed as limited to the examples set forth herein.

As discussed above, the market price paid to generation stations for supplying power to the grid often fluctuates due to various factors, including the need to maintain grid stability and based on current demand and usage by connected loads in distribution networks. Due to these factors, situations can arise where generation stations are offered substantially lower prices to deter an over-supply of power to the grid. Although these situations typically exist temporarily, generation stations are sometimes forced to either sell power to the grid at the much lower prices or adjust

8

operations to decrease the amount of power generated. Furthermore, some situations may even require generation stations to incur costs in order to offload power to the grid or to shut down generation temporarily.

The volatility in the market price offered for power supplied to the grid can be especially problematic for some types of generation stations. In particular, wind farms and some other types of renewable resource power producers may lack the ability to quickly adjust operations in response to changes in the market price offered for supplying power to the grid. As a result, power generation and management at some generation stations can be inefficient, which can frequently result in power being sold to the grid at low or negative prices. In some situations, a generation station may even opt to halt power generation temporarily to avoid such unfavorable pricing. As such, the time required to halt and to restart the power generation at a generation station can reduce the generation station's ability to take advantage of rising market prices for power supplied to the grid.

Example embodiments provided herein aim to assist generation stations in managing power generation operations and avoid unfavorable power pricing situations like those described above. In particular, example embodiments may involve providing a load that is positioned behind-the-meter ("BTM") and enabling the load to utilize power received behind-the-meter at a generation station in a timely manner. As a general rule of thumb, BTM power is not subject to traditional T&D costs.

For purposes herein, a generation station is considered to be configured for the primary purpose of generating utility-scale power for supply to the electrical grid (e.g., a Wide Area Synchronous Grid or a North American Interconnect).

In one embodiment, equipment located behind-the-meter ("BTM equipment") is equipment that is electrically connected to a generation station's power generation equipment behind (i.e., prior to) the generation station's POI with an electrical grid.

In one embodiment, behind-the-meter power ("BTM power") is electrical power produced by a generation station's power generation equipment and utilized behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM equipment receives power from the generation station, but the power received by the BTM equipment from the generation station has not passed through the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM power is utilized before being metered at the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

US 10,608,433 B1

9

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that supplies power to a grid, and the BTM equipment receives power from the generation station that is not subject to T&D charges, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that supplies power to a grid, and the BTM power is not subject to T&D charges before being used by electrical equipment, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, equipment may be considered behind-the-meter if the BTM equipment receives power generated from the generation station and that received power is not routed through the electrical grid before being delivered to the BTM equipment.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station, and BTM equipment receives that generated power, and that generated power received by the BTM equipment is not routed through the electrical grid before being delivered to the BTM equipment.

For purposes herein, BTM equipment may also be referred to as a behind-the-meter load ("BTM load") when the BTM equipment is actively consuming BTM power.

Beneficially, where BTM power is not subject to traditional T&D costs, a wind farm or other type of generation station can be connected to BTM loads which can allow the generation station to selectively avoid the adverse or less-than optimal cost structure occasionally associated with supplying power to the grid by shunting generated power to the BTM load.

An arrangement that positions and connects a BTM load to a generation station can offer several advantages. In such arrangements, the generation station may selectively choose whether to supply power to the grid or to the BTM load, or both. The operator of a BTM load may pay to utilize BTM power at a cost less than that charged through a consumer meter (e.g., 106d, 1060 located at a distribution network (e.g., 106a-c) receiving power from the grid. The operator of a BTM load may additionally or alternatively charge less than the market rate to consume excess power generated at the generation station during curtailment. As a result, the generation station may direct generated power based on the "best" price that the generation station can receive during a given time frame, and/or the lowest cost the generation station may incur from negative market pricing during curtailment. The "best" price may be the highest price that the generation station may receive for its generated power during a given duration, but can also differ within embodiments and may depend on various factors, such as a prior PPA.

In one example, by having a behind-the-meter option available, a generation station may transition from supplying all generated power to the grid to supplying some or all generated power to one or more BTM loads when the market price paid for power by grid operators drops below a predefined threshold (e.g., the price that the operator of the BTM load is willing to pay the generation station for power). Thus, by having an alternative option for power consumption (i.e., one or more BTM loads), the generation station can selectively utilize the different options to maximize the price received for generated power. In addition, the generation station may also utilize a BTM load to avoid or reduce

10

the economic impact in situations when supplying power to the grid would result in the generation station incurring a net cost.

Providing BTM power to a load can also benefit the BTM load operator. A BTM load may be able to receive and utilize BTM power received from the generation station at a cost that is lower than the cost for power from the grid (e.g., at a customer meter 106d, 1060. This is primarily due to the avoidance (or significant reduction) in T&D costs and the market effects of curtailment. As indicated above, the generation station may be willing to divert generated power to the BTM load rather than supplying the grid due to changing market conditions, or during maintenance periods, or for other non-market conditions. Thus, some situations may arise where the generation station offers power to the BTM load at a price that is substantially lower than the price available on the grid. Furthermore, in some situations, the BTM load may even be able to obtain and utilize BTM power from a generation station at no cost or even at negative pricing since the generation station may rather supply the BTM load with generated power during a given time range instead of paying a higher price for the grid to take the power or modifying operations to decrease power output.

Another example of cost-effective use of BTM power is when the generation station 202 is selling power to the grid at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price the generation station 202 would have to pay to the grid power to offload generation's station 202 generated power. Advantageously, one or more flexible datacenters 220 may take the generated power behind-the-meter, thereby allowing the generation station 202 to produce and obtain the production tax credit, while selling less power to the grid at the negative price.

Another example of cost-effective behind-the-meter power is when the generation station 202 is selling power to the grid at a negative price because the grid is oversupplied and/or the generation station 202 is instructed to stand down and stop producing altogether. A grid operator may select and direct certain generation stations to go offline and stop supplying power to the grid. Advantageously, one or more flexible datacenters may be used to take power behind-the-meter, thereby allowing the generation station 202 to stop supplying power to the grid, but still stay online and make productive use of the power generated.

Another example of beneficial behind-the-meter power use is when the generation station 202 is producing power that is, with reference to the grid, unstable, out of phase, or at the wrong frequency, or the grid is already unstable, out of phase, or at the wrong frequency. A grid operator may select certain generation stations to go either offline and stop producing power, or to take corrective action with respect to the grid power stability, phase, or frequency. Advantageously, one or more flexible datacenters 220 may be used to selectively consume power behind-the-meter, thereby allowing the generation station 202 to stop providing power to the grid and/or provide corrective feedback to the grid.

Another example of beneficial behind-the-meter power use is that cost-effective behind-the-meter power availability may occur when the generation station 202 is starting up or testing. Individual equipment in the power generation equipment 210 may be routinely offline for installation, maintenance, and/or service and the individual units must be tested prior to coming online as part of overall power generation equipment 210. During such testing or maintenance time, one or more flexible datacenters may be intermittently

US 10,608,433 B1

11

powered by the one or more units of the power generation equipment **210** that are offline from the overall power generation equipment **210**.

Another example of beneficial behind-the-meter power use is that datacenter control systems at the flexible datacenters **220** may quickly ramp up and ramp down power consumption by computing systems in the flexible datacenters **220** based on power availability from the generation station **202**. For instance, if the grid requires additional power and signals the demand via a higher local price for power, the generation station **202** can supply the grid with power nearly instantly by having active flexible datacenters **220** quickly ramp down and turn off computing systems (or switch to a stored energy source), thereby reducing an active BTM load.

Another example of beneficial behind-the-meter power use is in new photovoltaic generation stations **202**. For example, it is common to design and build new photovoltaic generation stations with a surplus of power capacity to account for degradation in efficiency of the photovoltaic panels over the life of the generation stations. Excess power availability at the generation station can occur when there is excess local power generation and/or low grid demand. In high incident sunlight situations, a photovoltaic generation station **202** may generate more power than the intended capacity of generation station **202**. In such situations, a photovoltaic generation station **202** may have to take steps to protect its equipment from damage, which may include taking one or more photovoltaic panels offline or shunting their voltage to dummy loads or the ground. Advantageously, one or more flexible datacenters (e.g., the flexible datacenters **220**) may take power behind-the-meter at the Generation Station **202**, thereby allowing the generation station **202** to operate the power generation equipment **210** within operating ranges while the flexible datacenters **220** receive BTM power without transmission or distribution costs.

Thus, for at least the reasons described herein, arrangements that involves providing a BTM load as an alternative option for a generation station to direct its generated power to can serve as a mutually beneficial relationship in which both the generation station and the BTM load can economically benefit. The above-noted examples of beneficial use of BTM power are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as benefits to unutilized BTM power capacity, BTM power pricing, or BTM power consumption.

Within example embodiments described herein, various types of utility-scale power producers may operate as generation stations **202** that are capable of supplying power to one or more loads behind-the-meter. For instance, renewable energy sources (e.g., wind, solar, hydroelectric, wave, water current, tidal), fossil fuel power generation sources (coal, natural gas), and other types of power producers (e.g., nuclear power) may be positioned in an arrangement that enables the intermittent supply of generated power behind-the-meter to one or more BTM loads. One of ordinary skill in the art will recognize that the generation station **202** may vary based on an application or design in accordance with one or more example embodiments.

In addition, the particular arrangement (e.g., connections) between the generation station and one or more BTM loads can vary within examples. In one embodiment, a generation station may be positioned in an arrangement wherein the generation station selectively supplies power to the grid and/or to one or more BTM loads. As such, power cost-analysis and other factors (e.g., predicted weather condi-

12

tions, contractual obligations, etc.) may be used by the generation station, a BTM load control system, a remote master control system, or some other system or enterprise, to selectively output power to either the grid or to one or more BTM loads in a manner that maximizes revenue to the generation station. In such an arrangement, the generation station may also be able to supply both the grid and one or more BTM loads simultaneously. In some instances, the arrangement may be configured to allow dynamic manipulation of the percentage of the overall generated power that is supplied to each option at a given time. For example, in some time periods, the generation station may supply no power to the BTM load.

In addition, the type of loads that are positioned behind-the-meter can vary within example embodiments. In general, a load that is behind-the-meter may correspond to any type of load capable of receiving and utilizing power behind-the-meter from a generation station. Some examples of loads include, but are not limited to, datacenters and electric vehicle (EV) charging stations.

Preferred BTM loads are loads that can be subject to intermittent power supply because BTM power may be available intermittently. In some instances, the generation station may generate power intermittently. For example, wind power station **102c** and/or photovoltaic power station **102d** may only generate power when resource are available or favorable. Additionally or alternatively, BTM power availability at a generation station may only be available intermittently due to power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and/or operational directives from grid operators or generation station operators.

Some example embodiments of BTM loads described herein involve using one or more computing systems to serve as a BTM load at a generation station. In particular, the computing system or computing systems may receive power behind-the-meter from the generation station to perform various computational operations, such as processing or storing information, performing calculations, mining for cryptocurrencies, supporting blockchain ledgers, and/or executing applications, etc.

Multiple computing systems positioned behind-the-meter may operate as part of a "flexible" datacenter that is configured to operate only intermittently and to receive and utilize BTM power to carry out various computational operations similar to a traditional datacenter. In particular, the flexible datacenter may include computing systems and other components (e.g., support infrastructure, a control system) configured to utilize BTM power from one or more generation stations. The flexible datacenter may be configured to use particular load ramping abilities (e.g., quickly increase or decrease power usage) to effectively operate during intermittent periods of time when power is available from a generation station and supplied to the flexible datacenter behind-the-meter, such as during situations when supplying generated power to the grid is not favorable for the generation station.

In some instances, the amount of power consumed by the computing systems at a flexible datacenter can be ramped up and down quickly, and potentially with high granularity (i.e., the load can be changed in small increments if desired). This may be done based on monitored power system conditions or other information analyses as discussed herein. As recited above, this can enable a generation station to avoid negative power market pricing and to respond quickly to grid direc-

US 10,608,433 B1

13

tives. And by extension, the flexible datacenter may obtain BTM power at a price lower than the cost for power from the grid.

Various types of computing systems can provide granular power ramping. Preferably, the computing systems can perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp down or up. In some embodiments, a control system may be used to activate or de-activate one or more computing systems in an array of computing systems. For example, the control system may provide control instructions to one or more blockchain miners (e.g., a group of blockchain miners), including instructions for powering on or off, adjusting frequency of computing systems performing operations (e.g., adjusting the processing frequency), adjusting the quantity of operations being performed, and when to operate within a low power mode (if available).

Within examples, a control system may correspond to a specialized computing system or may be a computing system within a datacenter serving in the role of the control system. The location of the control system can vary within examples as well. For instance, the control system may be located at a datacenter or physically separate from the datacenter. In some examples, the control system may be part of a network of control systems that manage computational operations, power consumption, and other aspects of a fleet of datacenters. The fleet of datacenters may include one or more traditional datacenters and/or flexible datacenters.

Some embodiments may involve using one or more control systems to direct time-insensitive (e.g., interruptible) computational tasks to computational hardware, such as central processing units (CPUs) and graphics processing units (GPUs), sited behind the meter, while other hardware is sited in front of the meter (i.e., consuming metered grid power via a customer meter (e.g., 106d, 1060) and possibly remote from the behind-the-meter hardware. As such, parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, machine learning, neural network processing, queued operations, and oil and gas field simulation models, are good candidates for such interruptible computational operations.

FIG. 2 shows a behind-the-meter arrangement with optional grid-power, including one or more flexible datacenters, according to one or more example embodiments. Dark arrows illustrate a typical power delivery direction. Consistent with FIG. 1, the arrangement illustrates a generation station 202 in the generation segment 102 of a Wide-Area Synchronous Grid. The generation station 202 supplies utility-scale power (typically >50 MW) via a generation power connection 250 to the Point of Interconnection 103 between the generation station 202 and the rest of the grid. Typically, the power supplied on connection 250 may be at 34.5 kV AC, but it may be higher or lower. Depending on the voltage at connection 250 and the voltage at transmission lines 104a, a transformer system 203 may step up the power supplied from the generation station 202 to high voltage (e.g., 115 kV+AC) for transmission over connection 252 and onto transmission lines 104a of transmission segment 104. Grid power carried on the transmission segment 104 may be from generation station 202 as well as other generation stations (not shown). Also consistent with FIG. 1, grid power is consumed at one or more distribution networks, including example distribution network 206. Grid power may be taken from the transmission lines 104a via connector 254 and stepped down to distribution network

14

voltages (e.g., typically 4 kV to 26 kV AC) and sent into the distribution networks, such as distribution network 206 via distribution line 256. The power on distribution line 256 may be further stepped down (not shown) before entering individual consumer facilities such as a remote master control system 262 and/or traditional datacenters 260 via customer meters 206A, which may correspond to customer meters 106d in FIG. 1, or customer meters 106f in FIG. 1 if the respective consumer facility includes a local customer power system, such as 106e (not shown in FIG. 2).

Consistent with FIG. 1, power entering the grid from generation station 202 is metered by a utility-scale generation-side meter. A utility-scale generation-side meter 253 is shown on the low side of transformer system 203 and an alternative location is shown as 253A on the high side of transformer system 203. Both locations may be considered settlement metering points for the generation station 202 at the POI 103. Alternatively, a utility-scale generation-side meter for the generation station 202 may be located at another location consistent with the descriptions of such meters provided herein.

Generation station 202 includes power generation equipment 210, which may include, as examples, wind turbines and/or photovoltaic panels. Power generation equipment 210 may further include other electrical equipment, including but not limited to switches, busses, collectors, inverters, and power unit transformers (e.g., transformers in wind turbines).

As illustrated in FIG. 2, generation station 202 is configured to connect with BTM equipment which may function as BTM loads. In the illustrated embodiment of FIG. 2, the BTM equipment includes flexible datacenters 220. Various configurations to supply BTM power to flexible datacenters 220 within the arrangement of FIG. 2 are described herein.

In one configuration, generated power may travel from the power generation equipment 210 over one or more connectors 230A, 230B to one or more electrical busses 240A, 240B, respectively. Each of the connectors 230A, 230B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 230B is shown with an open switch, and connector 230A is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used in various embodiments when BTM power is supplied without significant power conversion to BTM loads.

In various configurations, the busses 240A and 240B may be separated by an open switch 240C or combined into a common bus by a closed switch 240C.

In another configuration, generated power may travel from the power generation equipment 210 to the high side of a local step-down transformer 214. The generated power may then travel from the low side of the local step-down transformer 214 over one or more connectors 232A, 232B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 232A, 232B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 232A is shown with an open switch, and connector 232B is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the power generation equipment 210, but the generated power must be stepped down prior to use at the BTM loads.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of a local step-up transformer 212. The generated power may

15

then travel from the high side of the local step-up transformer 212 over one or more connectors 234A, 234B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 234A, 234B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 234A, 234B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of the local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 to the high side of local step-down transformer 213. The generated power may then travel from the low side of the local step-down transformer 213 over one or more connectors 236A, 236B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 236A, 236B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 236A, 236B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212, but the power must be stepped down prior to use at the BTM loads.

In one embodiment, power generated at the generation station 202 may be used to power a generation station control system 216 located at the generation station 202, when power is available. The generation station control system 216 may typically control the operation of the generation station 202. Generated power used at the generation station control system 216 may be supplied from bus 240A via connector 216A and/or from bus 240B via connector 216B. Each of the connectors 216A, 216B may be a switched connector such that power may be routed independently to 240A and/or 240B. While the generation station control system 216 can consume BTM power when powered via bus 240A or bus 240B, the BTM power taken by generation station control system 216 is insignificant in terms of rendering an economic benefit. Further, the generation station control system 216 is not configured to operate intermittently, as it generally must remain always on. Further still, the generation station control system 216 does not have the ability to quickly ramp a BTM load up or down.

In another embodiment, grid power may alternatively or additionally be used to power the generation station control system 216. As illustrated here, metered grid power from a distribution network, such as distribution network 206 for simplicity of illustration purposes only, may be used to power generation station control system 216 over connector 216C. Connector 216C may be a switched connector so that metered grid power to the generation station control system 216 can be switched on or off as needed. More commonly, metered grid power would be delivered to the generation station control system 216 via a separate distribution network (not shown), and also over a switched connector. Any such grid power delivered to the generation station control system 216 is metered by a customer meter 206A and subject to T&D costs.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202

16

through POI 103 and such grid power may power the generation station control system 216.

In some configurations, an energy storage system 218 may be connected to the generation station 202 via connector 218A, which may be a switched connector. For illustrative purposes only, connector 218A is shown with an open switch but in some embodiments it may be closed. The energy storage system 218 may be connected to bus 240A and/or bus 240B and store energy produced by the power generation equipment 210. The energy storage system may also be isolated from generation station 202 by switch 242A. In times of need, such as when the power generation equipment in an idle or off state and not generating power, the energy storage system may feed power to, for example, the flexible datacenters 220. The energy storage system may also be isolated from the flexible datacenters 220 by switch 242B.

In a preferred embodiment, as illustrated, power generation equipment 210 supplies BTM power via connector 242 to flexible datacenters 220. The BTM power used by the flexible datacenters 220 was generated by the generation station 202 and did not pass through the POI 103 or utility-scale generation-side meter 253, and is not subject to T&D charges. Power received at the flexible datacenters 220 may be received through respective power input connectors 220A. Each of the respective connectors 220A may be a switched connector that can electrically isolate the respective flexible datacenter 220 from the connector 242. Power equipment 220B may be arranged between the flexible datacenters 220 and the connector 242. The power equipment 220B may include, but is not limited to, power conditioners, unit transformers, inverters, and isolation equipment. As illustrated, each flexible datacenter 220 may be served by a respective power equipment 220B. However, in another embodiment, one power equipment 220B may serve multiple flexible datacenter 220.

In one embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter and electrically connected to the power generation equipment 210 behind (i.e., prior to) the generation station's POI 103 with the rest of the electrical grid.

In one embodiment, BTM power produced by the power generation equipment 210 is utilized by the flexible datacenters 220 behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as the flexible datacenters 220 are electrically connected to the generation station 202, and generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), and the flexible datacenters 220 receive power from the generation station 202, but the power received by the flexible datacenters 220 from the generation station 202 has not passed through a utility-scale generation-side meter. In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where power produced at the generation station 202 is subject to metering by utility-scale generation-side meter

253 (or 253A, or another utility-scale generation-side meter), but the BTM power supplied to the flexible datacenters 220 is utilized before being metered at the utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter). In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as they are electrically connected to the generation station 202 that supplies power to the grid, and the flexible datacenters 220 receive power from the generation station 202 that is not subject to T&D charges, but power otherwise received from the grid that is supplied by the generation station 202 is subject to T&D charges.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 that supplies power to a grid, and the generated power is not subject to T&D charges before being used by flexible datacenters 220, but power otherwise received from the connected grid is subject to T&D charges.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter because they receive power generated from the generation station 202 intended for the grid, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 for distribution to the grid, and the flexible datacenters 220 receive that power, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, metered grid power may alternatively or additionally be used to power one or more of the flexible datacenters 220, or a portion within one or more of the flexible datacenters 220. As illustrated here for simplicity, metered grid power from a distribution network, such as distribution network 206, may be used to power one or more flexible datacenters 220 over connector 256A and/or 256B. Each of connector 256A and/or 256B may be a switched connector so that metered grid power to the flexible datacenters 220 can be switched on or off as needed. More commonly, metered grid power would be delivered to the flexible datacenters 220 via a separate distribution network (not shown), and also over switched connectors. Any such grid power delivered to the flexible datacenters 220 is metered by customer meters 206A and subject to T&D costs. In one embodiment, connector 256B may supply metered grid power to a portion of one or more flexible datacenters 220. For example, connector 256B may supply metered grid power to control and/or communication systems for the flexible datacenters 220 that need constant power and cannot be subject to intermittent BTM power. Connector 242 may supply solely BTM power from the generation station 202 to high power demand computing systems within the flexible datacenters 220, in which case at least a portion of each flexible datacenters 220 so connected is operating as a BTM load. In another embodiment, connector 256A and/or 256B may supply all power used at one or more of the flexible

datacenters 220, in which case each of the flexible datacenters 220 so connected would not be operating as a BTM load.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the flexible datacenters 220.

The flexible datacenters 220 are shown in an example arrangement relative to the generation station 202. Particularly, generated power from the generation station 202 may be supplied to the flexible datacenters 220 through a series of connectors and/or busses (e.g., 232B, 240B, 242, 220A). As illustrated, in other embodiments, connectors between the power generation equipment 210 and other components may be switched open or closed, allowing other pathways for power transfer between the power generation equipment 210 and components, including the flexible datacenters 220. Additionally, the connector arrangement shown is illustrative only and other circuit arrangements are contemplated within the scope of supplying BTM power to a BTM load at generation station 202. For example, there may be more or fewer transformers, or one or more of transformers 212, 213, 214 may be transformer systems with multiple steppings and/or may include additional power equipment including but not limited to power conditioners, filters, switches, inverters, and/or AC/DC-DC/AC isolators. As another example, metered grid power connections to flexible datacenters 220 are shown via both 256A and 256B; however, a single connection may connect one or more flexible datacenters 220 (or power equipment 220B) to metered grid power and the one or more flexible datacenters 220 (or power equipment 220B) may include switching apparatus to direct BTM power and/or metered grid power to control systems, communication systems, and/or computing systems as desired.

In some examples, BTM power may arrive at the flexible datacenters 220 in a three-phase AC format. As such, power equipment (e.g., power equipment 220B) at one or more of the flexible datacenters 220 may enable each flexible datacenter 220 to use one or more phases of the power. For instance, the flexible datacenters 220 may utilize power equipment (e.g., power equipment 220B, or alternatively or additionally power equipment that is part of the flexible datacenter 220) to convert BTM power received from the generation station 202 for use at computing systems at each flexible datacenter 220. In other examples, the BTM power may arrive at one or more of the flexible datacenters 220 as DC power. As such, the flexible datacenters 220 may use the DC power to power computing systems. In some such examples, the DC power may be routed through a DC-to-DC converter that is part of power equipment 220B and/or flexible datacenter 220.

In some configurations, a flexible datacenter 220 may be arranged to only have access to power received behind-the-meter from a generation station 202. In the arrangement of FIG. 2, the flexible datacenters 220 may be arranged only with a connection to the generation station 202 and depend solely on power received behind-the-meter from the generation station 202. Alternatively or additionally, the flexible datacenters 220 may receive power from energy storage system 218.

In some configurations, one or more of the flexible datacenters 220 can be arranged to have connections to multiple sources that are capable of supplying power to a flexible datacenter 220. To illustrate a first example, the flexible datacenters 220 are shown connected to connector 242, which can be connected or disconnected via switches to

the energy storage system 218 via connector 218A, the generation station 202 via bus 240B, and grid power via metered connector 256A. In one embodiment, the flexible datacenters 220 may selectively use power received behind-the-meter from the generation station 202, stored power supplied by the energy storage system 218, and/or grid power. For instance, flexible datacenters 220 may use power stored in the energy storage system 218 when costs for using power supplied behind-the-meter from the generation station 202 are disadvantageous. By having access to the energy storage system 218 available, the flexible datacenters 220 may use the stored power and allow the generation station 202 to subsequently refill the energy storage system 218 when cost for power behind-the-meter is low. Alternatively, the flexible datacenters 220 may use power from multiple sources simultaneously to power different components (e.g., a first set and a second set of computing systems). Thus, the flexible datacenters 220 may leverage the multiple connections in a manner that can reduce the cost for power used by the computing systems at the flexible datacenters 220. The flexible datacenters 220 control system or the remote master control system 262 may monitor power conditions and other factors to determine whether the flexible datacenters 220 should use power from either the generation station 202, grid power, the energy storage system 218, none of the sources, or a subset of sources during a given time range. Other arrangements are possible as well. For example, the arrangement of FIG. 2 illustrates each flexible datacenter 220 as connected via a single connector 242 to energy storage system 218, generation station 202, and metered grid power via 256A. However, one or more flexible datacenters 220 may have independent switched connections to each energy source, allowing the one or more flexible datacenters 220 to operate from different energy sources than other flexible datacenters 220 at the same time.

The selection of which power source to use at a flexible datacenter (e.g., the flexible datacenters 220) or another type of BTM load can change based on various factors, such as the cost and availability of power from both sources, the type of computing systems using the power at the flexible datacenters 220 (e.g., some systems may require a reliable source of power for a long period), the nature of the computational operations being performed at the flexible datacenters 220 (e.g., a high priority task may require immediate completion regardless of cost), and temperature and weather conditions, among other possible factors. As such, a datacenter control system at the flexible datacenters 220, the remote master control system 262, or another entity (e.g., an operator at the generation station 202) may also influence and/or determine the source of power that the flexible datacenters 220 use at a given time to complete computational operations.

In some example embodiments, the flexible datacenters 220 may use power from the different sources to serve different purposes. For example, the flexible datacenters 220 may use metered power from grid power to power one or more systems at the flexible datacenters 220 that are configured to be always-on (or almost always on), such as a control and/or communication system and/or one or more computing systems (e.g., a set of computing systems performing highly important computational operations). The flexible datacenters 220 may use BTM power to power other components within the flexible datacenters 220, such as one or more computing systems that perform less critical computational operations.

In some examples, one or more flexible datacenters 220 may be deployed at the generation station 202. In other

examples, flexible datacenters 220 may be deployed at a location geographically remote from the generation station 202, while still maintaining a BTM power connection to the generation station 202.

In another example arrangement, the generation station 202 may be connected to a first BTM load (e.g., a flexible datacenter 220) and may supply power to additional BTM loads via connections between the first BTM load and the additional BTM loads (e.g., a connection between a flexible datacenter 220 and another flexible datacenter 220).

The arrangement in FIG. 2, and components included therein, are for non-limiting illustration purposes and other arrangements are contemplated in examples. For instance, in another example embodiment, the arrangement of FIG. 2 may include more or fewer components, such as more BTM loads, different connections between power sources and loads, and/or a different number of datacenters. In addition, some examples may involve one or more components within the arrangement of FIG. 2 being combined or further divided.

Within the arrangement of FIG. 2, a control system, such as the remote master control system 262 or another component (e.g., a control system associated with the grid operator, the generation station control system 216, or a datacenter control system associated with a traditional datacenter or one or more flexible datacenters) may use information to efficiently manage various operations of some of the components within the arrangement of FIG. 2. For example, the remote master control system 262 or another component may manage distribution and execution of computational operations at one or more traditional datacenters 260 and/or flexible datacenters 220 via one or more information-processing algorithms. These algorithms may utilize past and current information in real-time to manage operations of the different components. These algorithms may also make some predictions based on past trends and information analysis. In some examples, multiple computing systems may operate as a network to process information.

Information used to make decisions may include economic and/or power-related information, such as monitored power system conditions. Monitored power system conditions may include one or more of excess power generation at a generation station 202, excess power at a generation station 202 that a connected grid cannot receive, power generation at a generation station 202 subject to economic curtailment, power generation at a generation station 202 subject to reliability curtailment, power generation at a generation station 202 subject to power factor correction, low power generation at a generation station 202, start up conditions at a generation station 202, transient power generation conditions at a generation station 202, or testing conditions where there is an economic advantage to using behind-the-meter power generation at a generation station 202. These different monitored power system conditions can be weighted differently during processing and analysis.

In some examples, the information can include the cost for power from available sources (e.g., BTM power at the generation station 202 versus metered grid power) to enable comparisons to be made which power source costs less. In some instances, the information may include historic prices for power to enable the remote master control system 262 or another system to predict potential future prices in similar situations (e.g., the cost of power tends to trend upwards for grid power during warmer weather and peak-use hours). The information may also indicate the availability of power from the various sources (e.g., BTM power at the generation

US 10,608,433 B1

21

station 262, the energy storage system 218 at the generation station 262, and/or metered grid power).

In addition, the information may also include other data, including information associated with operations at components within the arrangement. For instance, the information may include data associated with performance of operations at the flexible datacenters 220 and the traditional datacenters 260, such as the number of computational tasks currently being performed, the types of tasks being performed (e.g., type of computational operation, time-sensitivity, etc.), the number, types, and capabilities of available computing systems, the amount of computational tasks awaiting performance, and the types of computing systems at one or more datacenters, among others. The information may also include data specifying the conditions at one or more datacenters (e.g., whether or not the temperatures are in a desired range, the amount of power available within an energy storage system such as 218), the amount of computational tasks awaiting performance in the queue of one or more of the datacenters, and the identities of the entities associated with the computational operations at one or more of the datacenters. Entities associated with computational operations may be, for example, owners of the datacenters, customers who purchase computational time at the datacenters, or other entities.

The information used by the remote master control system 262 or another component may include data associated with the computational operations to be performed, such as deadlines, priorities (e.g., high vs. low priority tasks), cost to perform based on required computing systems, the optimal computing systems (e.g., CPU vs GPU vs ASIC; processing unit capabilities, speeds, or frequencies, or instructional sets executable by the processing units) for performing each requested computational task, and prices each entity (e.g., company) is willing to pay for computational operations to be performed or otherwise supported via computing systems at a traditional datacenter 260 or a flexible datacenter 220, among others. In addition, the information may also include other data (e.g., weather conditions at locations of datacenters or power sources, any emergencies associated with a datacenter or power source, or the current value of bids associated with an auction for computational tasks).

The information may be updated in-real time and used to make the different operational decisions within the arrangement of FIG. 2. For instance, the information may help a component (e.g., the remote master control system 262 or a control system at a flexible datacenter 220) determine when to ramp up or ramp down power use at a flexible datacenter 220 or when to switch one or more computing systems at a flexible datacenter 220 into a low power mode or to operate at a different frequency, among other operational adjustments. The information can additionally or alternatively help a component within the arrangement of FIG. 2 to determine when to transfer computational operations between computing systems or between datacenters based on various factors. In some instances, the information may also be used to determine when to temporarily stop performing a computational operation or when to perform a computational operation at multiple sites for redundancy or other reasons. The information may further be used to determine when to accept new computational operations from entities or when to temporarily suspend accepting new tasks to be performed due to lack of computing system availability.

The remote master control system 262 represents a computing system that is capable of obtaining, managing, and using the information described above to manage and over-

22

see one or more operations within the arrangement of FIG. 2. As such, the remote master control system 262 may be one or more computing systems configured to process all, or a subset of, the information described above, such as power, environment, computational characterization, and economic factors to assist with the distribution and execution of computing operations among one or more datacenters. For instance, the remote master control system 262 may be configured to obtain and delegate computational operations among one or more datacenters based on a weighted analysis of a variety of factors, including one or more of the cost and availability of power, the types and availability of the computing systems at each datacenter, current and predicted weather conditions at the different locations of flexible datacenters (e.g., flexible datacenters 220) and generation stations (e.g., generation stations 202), levels of power storage available at one or more energy storage systems (e.g., energy storage system 218), and deadlines and other attributes associated with particular computational operations, among other possible factors. As such, the analysis of information performed by the remote master control system 262 may vary within examples. For instance, the remote master control system 262 may use real-time information to determine whether or not to route a computational operation to a particular flexible datacenter (e.g., a flexible datacenter 220) or to transition a computational operation between datacenters (e.g., from traditional datacenter 260 to a flexible datacenter 220).

As shown in FIG. 2, the generation station 202 may be able to supply power to the grid and/or BTM loads such as flexible datacenters 220. With such a configuration, the generation station 202 may selectively provide power to the BTM loads and/or the grid based on economic and power availability considerations. For example, the generation station 202 may supply power to the grid when the price paid for the power exceeds a particular threshold (e.g., the power price offered by operators of the flexible datacenters 220). In some instances, the operator of a flexible datacenter and the operator of a generation station capable of supplying BTM power to the flexible datacenter may utilize a predefined arrangement (e.g., a contract) that specifies a duration and/or price range when the generation station may supply power to the flexible datacenter.

The remote master control system 262 may be capable of directing one or more flexible datacenters 220 to ramp-up or ramp-down to desired power consumption levels, and/or to control cooperative action of multiple flexible datacenters by determining how to power each individual flexible datacenter 220 in accordance with operational directives.

The configuration of the remote master control system 262 can vary within examples as further discussed with respect to FIGS. 2, 3, and 7-9. The remote master control system 262 may operate as a single computing system or may involve a network of computing systems. Preferably, the remote master control system 262 is implemented across one or more servers in a fault-tolerant operating environment that ensures continuous uptime and connectivity by virtue of its distributed nature. Alternatively, although the remote master control system 262 is shown as a physically separate component arrangement for FIG. 2, the remote master control system 262 may be combined with another component in other embodiments. To illustrate an example, the remote master control system 262 may operate as part of a flexible datacenter (e.g., a computing system or a datacenter control system of the flexible datacenter 220), includ-

US 10,608,433 B1

23

ing sharing components with a flexible datacenter, sharing power with a flexible datacenter, and/or being co-located with a flexible datacenter.

In addition, the remote master control system 262 may communicate with components within the arrangement of FIG. 2 using various communication technologies, including wired and wireless communication technologies. For instance, the remote master control system 262 may use wired (not illustrated) or wireless communication to communicate with datacenter control systems or other computing systems at the flexible datacenters 220 and the traditional datacenters 260. The remote master control system 262 may also communicate with entities inside or outside the arrangement of FIG. 2 and other components within the arrangement of FIG. 2 via wired or wireless communication. For instance, the remote master control system 262 may use wireless communication to obtain computational operations from entities seeking support for the computational operations at one or more datacenters in exchange for payment. The remote master control system 262 may communicate directly with the entities or may obtain the computational operations from the traditional datacenters 260. For instance, an entity may submit jobs (e.g., computational operations) to one or more traditional datacenters 260. The remote master control system 262 may determine that transferring one or more of the computational operations to a flexible datacenter 220 may better support the transferred computational operations. For example, the remote master control system 262 may determine that the transfer may enable the computational operations to be completed quicker and/or at a lower cost. In some examples, the remote master control system 262 may communicate with the entity to obtain approval prior to transferring the one or more computational operations.

The remote master control system 262 may also communicate with grid operators and/or an operator of generation station 202 to help determine power management strategies when distributing computational operations across the various datacenters. In addition, the remote master control system 262 may communicate with other sources, such as weather prediction systems, historical and current power price databases, and auction systems, etc.

In further examples, the remote master control system 262 or another computing system within the arrangement of FIG. 2 may use wired or wireless communication to submit bids within an auction that involves a bidder (e.g., the highest bid) obtaining computational operations or other tasks to be performed. Particularly, the remote master control system 262 may use the information discussed above to develop bids to obtain computing operations for performance at available computing systems at flexible datacenters (e.g., flexible datacenters 220).

In the example arrangement shown in FIG. 2, the flexible datacenters 220 represent example loads that can receive power behind-the-meter from the generation station 202. In such a configuration, the flexible datacenters 220 may obtain and utilize power behind-the-meter from the generation station 202 to perform various computational operations. Performance of a computational operation may involve one or more computing systems providing resources useful in the computational operation. For instance, the flexible datacenters 220 may include one or more computing systems configured to store information, perform calculations and/or parallel processes, perform simulations, mine cryptocurrencies, and execute applications, among other potential tasks. The computing systems can be specialized or generic and can be arranged at each flexible datacenter 220 in a variety

24

of ways (e.g., straight configuration, zig-zag configuration) as further discussed with respect to FIGS. 6A, 6B. Furthermore, although the example arrangement illustrated in FIG. 2 shows configurations where flexible datacenters 220 serve as BTM loads, other types of loads can be used as BTM loads within examples.

The arrangement of FIG. 2 includes the traditional datacenters 260 coupled to metered grid power. The traditional datacenters 260 using metered grid power to provide computational resources to support computational operations. One or more enterprises may assign computational operations to the traditional datacenters 260 with expectations that the datacenters reliably provide resources without interruption (i.e., non-intermittently) to support the computational operations, such as processing abilities, networking, and/or volatile storage. Similarly, one or more enterprises may also request computational operations to be performed by the flexible datacenters 220. The flexible datacenters 220 differ from the traditional datacenters 260 in that the flexible datacenters 220 are arranged and/or configured to be connected to BTM power, are expected to operate intermittently, and are expected to ramp load (and thus computational capability) up or down regularly in response to control directives. In some examples, the flexible datacenters 220 and the traditional datacenters 260 may have similar configurations and may only differ based on the source(s) of power relied upon to power internal computing systems. Preferably, however, the flexible datacenters 220 include particular fast load ramping abilities (e.g., quickly increase or decrease power usage) and are intended and designed to effectively operate during intermittent periods of time.

FIG. 3 shows a block diagram of the remote master control system 300 according to one or more example embodiments. Remote master control system 262 may take the form of remote master control system 300, or may include less than all components in remote master control system 300, different components than in remote master control system 300, and/or more components than in remote master control system 300.

The remote master control system 300 may perform one or more operations described herein and may include a processor 302, a data storage unit 304, a communication interface 306, a user interface 308, an operations and environment analysis module 310, and a queue system 312. In other examples, the remote master control system 300 may include more or fewer components in other possible arrangements.

As shown in FIG. 3, the various components of the remote master control system 300 can be connected via one or more connection mechanisms (e.g., a connection mechanism 314). In this disclosure, the term "connection mechanism" means a mechanism that facilitates communication between two or more devices, systems, components, or other entities. For instance, a connection mechanism can be a simple mechanism, such as a cable, PCB trace, or system bus, or a relatively complex mechanism, such as a packet-based communication network (e.g., LAN, WAN, and/or the Internet). In some instances, a connection mechanism can include a non-tangible medium (e.g., where the connection is wireless).

As part of the arrangement of FIG. 2, the remote master control system 300 (corresponding to remote master control system 262) may perform a variety of operations, such as management and distribution of computational operations among datacenters, monitoring operational, economic, and environment conditions, and power management. For instance, the remote master control system 300 may obtain

US 10,608,433 B1

25

computational operations from one or more enterprises for performance at one or more datacenters. The remote master control system 300 may subsequently use information to distribute and assign the computational operations to one or more datacenters (e.g., the flexible datacenters 220) that have the resources (e.g., particular types of computing systems and available power) available to complete the computational operations. In some examples, the remote master control system 300 may assign all incoming computational operation requests to the queue system 312 and subsequently assign the queued requests to computing systems based on an analysis of current market and power conditions.

Although the remote master control system 300 is shown as a single entity, a network of computing systems may perform the operations of the remote master control system 300 in some examples. For example, the remote master control system 300 may exist in the form of computing systems (e.g., datacenter control systems) distributed across multiple datacenters.

The remote master control system 300 may include one or more processors 302. As such, the processor 302 may represent one or more general-purpose processors (e.g., a microprocessor) and/or one or more special-purpose processors (e.g., a digital signal processor (DSP)). In some examples, the processor 302 may include a combination of processors within examples. The processor 302 may perform operations, including processing data received from the other components within the arrangement of FIG. 2 and data obtained from external sources, including information such as weather forecasting systems, power market price systems, and other types of sources or databases.

The data storage unit 304 may include one or more volatile, non-volatile, removable, and/or non-removable storage components, such as magnetic, optical, or flash storage, and/or can be integrated in whole or in part with the processor 302. As such, the data storage unit 304 may take the form of a non-transitory computer-readable storage medium, having stored thereon program instructions (e.g., compiled or non-compiled program logic and/or machine code) that, when executed by the processor 302, cause the remote master control system 300 to perform one or more acts and/or functions, such as those described in this disclosure. Such program instructions can define and/or be part of a discrete software application. In some instances, the remote master control system 300 can execute program instructions in response to receiving an input, such as from the communication interface 306, the user interface 308, or the operations and environment analysis module 310. The data storage unit 304 may also store other information, such as those types described in this disclosure.

In some examples, the data storage unit 304 may serve as storage for information obtained from one or more external sources. For example, data storage unit 304 may store information obtained from one or more of the traditional datacenters 260, a generation station 202, a system associated with the grid, and flexible datacenters 220. As examples only, data storage 304 may include, in whole or in part, local storage, dedicated server-managed storage, network attached storage, and/or cloud-based storage, and/or combinations thereof.

The communication interface 306 can allow the remote master control system 300 to connect to and/or communicate with another component according to one or more protocols. For instance, the communication interface 306 may be used to obtain information related to current, future, and past prices for power, power availability, current and predicted

26

weather conditions, and information regarding the different datacenters (e.g., current workloads at datacenters, types of computing systems available within datacenters, price to obtain power at each datacenter, levels of power storage available and accessible at each datacenter, etc.). In an example, the communication interface 306 can include a wired interface, such as an Ethernet interface or a high-definition serial-digital-interface (HD-SDI). In another example, the communication interface 406 can include a wireless interface, such as a cellular, satellite, WiMAX, or WI-FI interface. A connection can be a direct connection or an indirect connection, the latter being a connection that passes through and/or traverses one or more components, such as such as a router, switcher, or other network device. Likewise, a wireless transmission can be a direct transmission or an indirect transmission. The communication interface 306 may also utilize other types of wireless communication to enable communication with datacenters positioned at various locations.

The communication interface 306 may enable the remote master control system 300 to communicate with the components of the arrangement of FIG. 2. In addition, the communication interface 306 may also be used to communicate with the various datacenters, power sources, and different enterprises submitting computational operations for the datacenters to support.

The user interface 308 can facilitate interaction between the remote master control system 300 and an administrator or user, if applicable. As such, the user interface 308 can include input components such as a keyboard, a keypad, a mouse, a touch-sensitive panel, a microphone, and/or a camera, and/or output components such as a display device (which, for example, can be combined with a touch-sensitive panel), a sound speaker, and/or a haptic feedback system. More generally, the user interface 308 can include hardware and/or software components that facilitate interaction between remote master control system 300 and the user of the system.

In some examples, the user interface 308 may enable the manual examination and/or manipulation of components within the arrangement of FIG. 2. For instance, an administrator or user may use the user interface 308 to check the status of, or change, one or more computational operations, the performance or power consumption at one or more datacenters, the number of tasks remaining within the queue system 312, and other operations. As such, the user interface 308 may provide remote connectivity to one or more systems within the arrangement of FIG. 2.

The operations and environment analysis module 310 represents a component of the remote master control system 300 associated with obtaining and analyzing information to develop instructions/directives for components within the arrangement of FIG. 2. The information analyzed by the operations and environment analysis module 310 can vary within examples and may include the information described above with respect predicting and/or directing the use of BTM power. For instance, the operations and environment analysis module 310 may obtain and access information related to the current power state of computing systems operating as part of the flexible datacenters 220 and other datacenters that the remote master control system 300 has access to. This information may be used to determine when to adjust power usage or mode of one or more computing systems. In addition, the remote master control system 300 may provide instructions a flexible datacenter 220 to cause a subset of the computing systems to transition into a low power mode to consume less power while still performing

operations at a slower rate. The remote master control system 300 may also use power state information to cause a set of computing systems at a flexible datacenter 220 to operate at a higher power consumption mode. In addition, the remote master control system 300 may transition computing systems into sleep states or power on/off based on information analyzed by the operations and environment analysis module 310.

In some examples, the operations and environment analysis module 310 may use location, weather, activity levels at the flexible datacenters or the generation station, and power cost information to determine control strategies for one or more components in the arrangement of FIG. 2. For instance, the remote master control system 300 may use location information for one or more datacenters to anticipate potential weather conditions that could impact access to power. In addition, the operations and environment analysis module 310 may assist the remote master control system 300 determine whether to transfer computational operations between datacenters based on various economic and power factors.

The queue system 312 represents a queue capable of organizing computational operations to be performed by one or more datacenters. Upon receiving a request to perform a computational operation, the remote master control system 300 may assign the computational operation to the queue until one or more computing systems are available to support the computational operation. The queue system 312 may be used for organizing and transferring computational tasks in real time.

The organizational design of the queue system 312 may vary within examples. In some examples, the queue system 312 may organize indications (e.g., tags, pointers) to sets of computational operations requested by various enterprises. The queue system 312 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 312 may include one or more queues that operate using the FIFO data structure.

In some examples, one or more queues within the queue system 312 may use other designs of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others. Examples using a queue system are further described below with respect to FIG. 9.

In some examples, the remote master control system 300 may be configured to monitor one or more auctions to obtain computational operations for datacenters to support. Particularly, the remote master control system 300 may use resource availability and power prices to develop and submit bids to an external or internal auction system for the right to support particular computational operations. As a result, the remote master control system 300 may identify computational operations that could be supported at one or more flexible datacenters 220 at low costs.

FIG. 4 is a block diagram of a generation station 400, according to one or more example embodiments. Generation station 202 may take the form of generation station 400, or may include less than all components in generation station 400, different components than in generation station 400, and/or more components than in generation station 400. The generation station 400 includes power generation equipment

401, a communication interface 408, a behind-the-meter interface 406, a grid interface 404, a user interface 410, a generation station control system 414, and power transformation equipment 402. The power generation equipment 210 may take the form of power generation equipment 401, or may include less than all components in power generation equipment 401, different components than in power generation equipment 401, and/or more components than in power generation equipment 401. Generation station control system 216 may take the form of generation station control system 414, or may include less than all components in generation station control system 414, different components than in generation station control system 414, and/or more components than in generation station control system 414. Some or all of the components generation station 400 may be connected via a communication interface 516. These components are illustrated in FIG. 4 to convey an example configuration for the generation station 400 (corresponding to generation station 202 shown in FIG. 2). In other examples, the generation station 400 may include more or fewer components in other arrangements.

The generation station 400 can correspond to any type of grid-connected utility-scale power producer capable of supplying power to one or more loads. The size, amount of power generated, and other characteristics of the generation station 400 may differ within examples. For instance, the generation station 400 may be a power producer that provides power intermittently. The power generation may depend on monitored power conditions, such as weather at the location of the generation station 400 and other possible conditions. As such, the generation station 400 may be a temporary arrangement, or a permanent facility, configured to supply power. The generation station 400 may supply BTM power to one or more loads and supply metered power to the electrical grid. Particularly, the generation station 400 may supply power to the grid as shown in the arrangement of FIG. 2.

The power generation equipment 401 represents the component or components configured to generate utility-scale power. As such, the power generation equipment 401 may depend on the type of facility that the generation station 400 corresponds to. For instance, the power generation equipment 401 may correspond to electric generators that transform kinetic energy into electricity. The power generation equipment 401 may use electromagnetic induction to generate power. In other examples, the power generation equipment 401 may utilize electrochemistry to transform chemical energy into power. The power generation equipment 401 may use the photovoltaic effect to transform light into electrical energy. In some examples, the power generation equipment 401 may use turbines to generate power. The turbines may be driven by, for example, wind, water, steam or burning gas. Other examples of power production are possible.

The communication interface 408 can enable the generation station 400 to communicate with other components within the arrangement of FIG. 2. As such, the communication interface 408 may operate similarly to the communication interface 306 of the remote master control system 300 and the communication interface 503 of the flexible datacenter 500.

The generation station control system 414 may be one or more computing systems configured to control various aspects of the generation station 400.

The BTM interface 406 is a module configured to enable the power generation equipment 401 to supply BTM power to one or more loads and may include multiple components.

US 10,608,433 B1

29

The arrangement of the BTM interface 406 may differ within examples based on various factors, such as the number of flexible datacenters 220 (or 500) coupled to the generation station 400, the proximity of the flexible datacenters 220 (or 500), and the type of generation station 400, among others. In some examples, the BTM interface 406 may be configured to enable power delivery to one or more flexible datacenters positioned near the generation station 400. Alternatively, the BTM interface 406 may also be configured to enable power delivery to one or more flexible datacenters 220 (or 500) positioned remotely from the generation station 400.

The grid interface 404 is a module configured to enable the power generation equipment 401 to supply power to the grid and may include multiple components. As such, the grid interface 404 may couple to one or more transmission lines (e.g., transmission lines 404a shown in FIG. 2) to enable delivery of power to the grid.

The user interface 410 represents an interface that enables administrators and/or other entities to communicate with the generation station 400. As such, the user interface 410 may have a configuration that resembles the configuration of the user interface 308 shown in FIG. 3. An operator may utilize the user interface 410 to control or monitor operations at the generation station 400.

The power transformation equipment 402 represents equipment that can be utilized to enable power delivery from the power generation equipment 401 to the loads and to transmission lines linked to the grid. Example power transformation equipment 402 includes, but is not limited to, transformers, inverters, phase converters, and power conditioners.

FIG. 5 shows a block diagram of a flexible datacenter 500, according to one or more example embodiments. Flexible datacenters 220 may take the form of flexible datacenter 500, or may include less than all components in flexible datacenter 500, different components than in flexible datacenter 500, and/or more components than in flexible datacenter 500. In the example embodiment shown in FIG. 5, the flexible datacenter 500 includes a power input system 502, a communication interface 503, a datacenter control system 504, a power distribution system 506, a climate control system 508, one or more sets of computing systems 512, and a queue system 514. These components are shown connected by a communication bus 528. In other embodiments, the configuration of flexible datacenter 500 can differ, including more or fewer components. In addition, the components within flexible datacenter 500 may be combined or further divided into additional components within other embodiments.

The example configuration shown in FIG. 5 represents one possible configuration for a flexible datacenter. As such, each flexible datacenter may have a different configuration when implemented based on a variety of factors that may influence its design, such as location and temperature that the location, particular uses for the flexible datacenter, source of power supplying computing systems within the flexible datacenter, design influence from an entity (or entities) that implements the flexible datacenter, and space available for the flexible datacenter. Thus, the embodiment of flexible datacenter 220 shown in FIG. 2 represents one possible configuration for a flexible datacenter out of many other possible configurations.

The flexible datacenter 500 may include a design that allows for temporary and/or rapid deployment, setup, and start time for supporting computational operations. For instance, the flexible datacenter 500 may be rapidly

30

deployed at a location near a source of generation station power (e.g., near a wind farm or solar farm). Rapid deployment may involve positioning the flexible datacenter 500 at a target location and installing and/or configuring one or more racks of computing systems within. The racks may include wheels to enable swift movement of the computing systems. Although the flexible datacenter 500 could theoretically be placed anywhere, transmission losses may be minimized by locating it proximate to BTM power generation.

The physical construction and layout of the flexible datacenter 500 can vary. In some instances, the flexible datacenter 500 may utilize a metal container (e.g., a metal container 602 shown in FIG. 6A). In general, the flexible datacenter 500 may utilize some form of secure weather-proof housing designed to protect interior components from wind, weather, and intrusion. The physical construction and layout of example flexible datacenters are further described with respect to FIGS. 6A-6B.

Within the flexible datacenter 500, various internal components enable the flexible datacenter 500 to utilize power to perform some form of operations. The power input system 502 is a module of the flexible datacenter 500 configured to receive external power and input the power to the different components via assistance from the power distribution system 506. As discussed with respect to FIG. 2, the sources of external power feeding a flexible datacenter can vary in both quantity and type (e.g., the generation stations 202, 400, grid-power, energy storage systems). Power input system 502 includes a BTM power input sub-system 522, and may additionally include other power input sub-systems (e.g., a grid-power input sub-system 524 and/or an energy storage input sub-system 526). In some instances, the quantity of power input sub-systems may depend on the size of the flexible datacenter and the number and/or type of computing systems being powered. In an example embodiment, the flexible datacenter may use grid power as the primary power supply.

In some embodiments, the power input system 502 may include some or all of flexible datacenter Power Equipment 220B. The power input system 502 may be designed to obtain power in different forms (e.g., single phase or three-phase behind-the-meter alternating current ("AC") voltage, and/or direct current ("DC") voltage). As shown, the power input system 502 includes a BTM power input sub-system 522, a grid power input sub-system 524, and an energy input sub-system 526. These sub-systems are included to illustrate example power input sub-systems that the flexible datacenter 500 may utilize, but other examples are possible. In addition, in some instances, these sub-systems may be used simultaneously to supply power to components of the flexible datacenter 500. The sub-systems may also be used based on available power sources.

In some implementations, the BTM power input sub-system 522 may include one or more AC-to-AC step-down transformers used to step down supplied medium-voltage AC to low voltage AC (e.g., 120V to 600V nominal) used to power computing systems 512 and/or other components of flexible datacenter 500. The power input system 502 may also directly receive single-phase low voltage AC from a generation station as BTM power, from grid power, or from a stored energy system such as energy storage system 218. In some implementations, the power input system 502 may provide single-phase AC voltage to the datacenter control system 504 (and/or other components of flexible datacenter 500) independent of power supplied to computing systems 512 to enable the datacenter control system 504 to perform

management operations for the flexible datacenter 500. For instance, the grid power input sub-system 524 may use grid power to supply power to the datacenter control system 504 to ensure that the datacenter control system 504 can perform control operations and communicate with the remote master control system 300 (or 262) during situations when BTM power is not available. As such, the datacenter control system 504 may utilize power received from the power input system 502 to remain powered to control the operation of flexible datacenter 500, even if the computational operations performed by the computing system 512 are powered intermittently. In some instances, the datacenter control system 504 may switch into a lower power mode to utilize less power while still maintaining the ability to perform some functions.

The power distribution system 506 may distribute incoming power to the various components of the flexible datacenter 500. For instance, the power distribution system 506 may direct power (e.g., single-phase or three-phase AC) to one or more components within flexible datacenter 500. In some embodiments, the power distribution system 506 may include some or all of flexible datacenter Power Equipment 220B.

In some examples, the power input system 502 may provide three phases of three-phase AC voltage to the power distribution system 506. The power distribution system 506 may controllably provide a single phase of AC voltage to each computing system or groups of computing systems 512 disposed within the flexible datacenter 500. The datacenter control system 504 may controllably select which phase of three-phase nominal AC voltage that power distribution system 506 provides to each computing system 512 or groups of computing systems 512. This is one example manner in which the datacenter control system 504 may modulate power delivery (and load at the flexible datacenter 500) by ramping-up flexible datacenter 500 to fully operational status, ramping-down flexible datacenter 500 to offline status (where only datacenter control system 504 remains powered), reducing load by withdrawing power delivery from, or reducing power to, one or more of the computing systems 512 or groups of the computing systems 512, or modulating power factor correction for the generation station 300 (or 202) by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more of the computing systems 512 or groups of the computing systems 512. The datacenter control system 504 may direct power to certain sets of computing systems based on computational operations waiting for computational resources within the queue system 514. In some embodiments, the flexible datacenter 500 may receive BTM DC power to power the computing systems 512.

One of ordinary skill in the art will recognize that a voltage level of three-phase AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase AC voltage, single-phase AC voltage, and DC voltage may vary based on the application or design in accordance with one or more embodiments.

As discussed above, the datacenter control system 504 may perform operations described herein, such as dynamically modulating power delivery to one or more of the computing systems 512 disposed within flexible datacenter 500. For instance, the datacenter control system 504 may modulate power delivery to one or more of the computing

systems 512 based on various factors, such as BTM power availability or an operational directive from a generation station 262 or 300 control system, a remote master control system 262 or 300, or a grid operator. In some examples, the datacenter control system 504 may provide computational operations to sets of computing systems 512 and modulate power delivery based on priorities assigned to the computational operations. For instance, an important computational operation (e.g., based on a deadline for execution and/or price paid by an entity) may be assigned to a particular computing system or set of computing systems 512 that has the capacity, computational abilities to support the computational operation. In addition, the datacenter control system 504 may also prioritize power delivery to the computing system or set of computing systems 512.

In some example, the datacenter control system 504 may further provide directives to one or more computing systems to change operations in some manner. For instance, the datacenter control system 504 may cause one or more computing systems 512 to operate at a lower or higher frequency, change clock cycles, or operate in a different power consumption mode (e.g., a low power mode). These abilities may vary depending on types of computing systems 512 available at the flexible datacenter 500. As a result, the datacenter control system 504 may be configured to analyze the computing systems 512 available either on a periodic basis (e.g., during initial set up of the flexible datacenter 500) or in another manner (e.g., when a new computational operation is assigned to the flexible datacenter 500).

The datacenter control system 504 may also implement directives received from the remote master control system 262 or 300. For instance, the remote master control system 262 or 300 may direct the flexible datacenter 500 to switch into a low power mode. As a result, one or more of the computing systems 512 and other components may switch to the low power mode in response.

The datacenter control system 504 may utilize the communication interface 503 to communicate with the remote master control system 262 or 300, other datacenter control systems of other datacenters, and other entities. As such, the communication interface 503 may include components and operate similar to the communication interface 306 of the remote master control system 300 described with respect to FIG. 4.

The flexible datacenter 500 may also include a climate control system 508 to maintain computing systems 512 within a desired operational temperature range. The climate control system 508 may include various components, such as one or more air intake components, an evaporative cooling system, one or more fans, an immersive cooling system, an air conditioning or refrigerant cooling system, and one or more air outtake components. One of ordinary skill in the art will recognize that any suitable heat extraction system configured to maintain the operation of computing systems 512 within the desired operational temperature range may be used.

The flexible datacenter 500 may further include an energy storage system 510. The energy storage system 510 may store energy for subsequent use by computing systems 512 and other components of flexible datacenter 500. For instance, the energy storage system 510 may include a battery system. The battery system may be configured to convert AC voltage to DC voltage and store power in one or more storage cells. In some instances, the battery system may include a DC-to-AC inverter configured to convert DC

US 10,608,433 B1

33

voltage to AC voltage, and may further include an AC phase-converter, to provide AC voltage for use by flexible datacenter 500.

The energy storage system 510 may be configured to serve as a backup source of power for the flexible datacenter 500. For instance, the energy storage system 510 may receive and retain power from a BTM power source at a low cost (or no cost at all). This low-cost power can then be used by the flexible datacenter 500 at a subsequent point, such as when BTM power costs more. Similarly, the energy storage system 510 may also store energy from other sources (e.g., grid power). As such, the energy storage system 510 may be configured to use one or more of the sub-systems of the power input system 502.

In some examples, the energy storage system 510 may be external to the flexible datacenter 500. For instance, the energy storage system 510 may be an external source that multiple flexible datacenters utilize for back-up power.

The computing systems 512 represent various types of computing systems configured to perform computational operations. Performance of computational operations include a variety of tasks that one or more computing systems may perform, such as data storage, calculations, application processing, parallel processing, data manipulation, cryptocurrency mining, and maintenance of a distributed ledger, among others. As shown in FIG. 5, the computing systems 512 may include one or more CPUs 516, one or more GPUs 518, and/or one or more Application-Specific Integrated Circuits (ASIC's) 520. Each type of computing system 512 may be configured to perform particular operations or types of operations.

Due to different performance features and abilities associated with the different types of computing systems, the datacenter control system 504 may determine, maintain, and/or relay this information about the types and/or abilities of the computing systems, quantity of each type, and availability to the remote master control system 262 or 300 on a routine basis (e.g., periodically or on-demand). This way, the remote master control system 262 or 300 may have current information about the abilities of the computing systems 512 when distributing computational operations for performance at one or more flexible datacenters. Particularly, the remote master control system 262 or 300 may assign computational operations based on various factors, such as the types of computing systems available and the type of computing systems required by each computing operation, the availability of the computing systems, whether computing systems can operate in a low power mode, and/or power consumption and/or costs associated with operating the computing systems, among others.

The quantity and arrangement of these computing systems 512 may vary within examples. In some examples, the configuration and quantity of computing systems 512 may depend on various factors, such as the computational tasks that are performed by the flexible datacenter 500. In other examples, the computing systems 512 may include other types of computing systems as well, such as DSPs, SIMDs, neural processors, and/or quantum processors.

As indicated above, the computing systems 512 can perform various computational operations, including in different configurations. For instance, each computing system may perform a particular computational operation unrelated to the operations performed at other computing systems. Groups of the computing systems 512 may also be used to work together to perform computational operations.

In some examples, multiple computing systems may perform the same computational operation in a redundant

34

configuration. This redundant configuration creates a backup that prevents losing progress on the computational operation in situations of a computing failure or intermittent operation of one or more computing systems. In addition, the computing systems 512 may also perform computational operations using a check point system. The check point system may enable a first computing system to perform operations up to a certain point (e.g., a checkpoint) and switch to a second computing system to continue performing the operations from that certain point. The check point system may also enable the datacenter control system 504 to communicate statuses of computational operations to the remote master control system 262 or 300. This can further enable the remote master control system 262 or 300 to transfer computational operations between different flexible datacenters allowing computing systems at the different flexible datacenters to resume support of computational operations based on the check points.

The queue system 514 may operate similar to the queue system 312 of the remote master control system 300 shown in FIG. 3. Particularly, the queue system 514 may help store and organize computational tasks assigned for performance at the flexible datacenter 500. In some examples, the queue system 514 may be part of a distributed queue system such that each flexible datacenter in a fleet of flexible datacenter includes a queue, and each queue system 514 may be able to communicate with other queue systems. In addition, the remote master control system 262 or 300 may be configured to assign computational tasks to the queues located at each flexible datacenter (e.g., the queue system 514 of the flexible datacenter 500). As such, communication between the remote master control system 262 or 300 and the datacenter control system 504 and/or the queue system 514 may allow organization of computational operations for the flexible datacenter 500 to support.

FIG. 6A shows another structural arrangement for a flexible datacenter, according to one or more example embodiments. The particular structural arrangement shown in FIG. 6A may be implemented at flexible datacenter 500. The illustration depicts the flexible datacenter 500 as a mobile container 702 equipped with the power input system 502, the power distribution system 506, the climate control system 508, the datacenter control system 504, and the computing systems 512 arranged on one or more racks 604. These components of flexible datacenter 500 may be arranged and organized according to an example structural region arrangement. As such, the example illustration represents one possible configuration for the flexible datacenter 500, but others are possible within examples.

As discussed above, the structural arrangement of the flexible datacenter 500 may depend on various factors, such as the ability to maintain temperature within the mobile container 602 within a desired temperature range. The desired temperature range may depend on the geographical location of the mobile container 602 and the type and quantity of the computing systems 512 operating within the flexible datacenter 500 as well as other possible factors. As such, the different design elements of the mobile container 602 including the inner contents and positioning of components may depend on factors that aim to maximize the use of space within mobile container 602, lower the amount of power required to cool the computing systems 512, and make setup of the flexible datacenter 500 efficient. For instance, a first flexible datacenter positioned in a cooler geographic region may include less cooling equipment than a second flexible datacenter positioned in a warmer geographic region.

US 10,608,433 B1

35

As shown in FIG. 6A, the mobile container 602 may be a storage trailer disposed on permanent or removable wheels and configured for rapid deployment. In other embodiments, the mobile container 602 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, the mobile container 602 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile flexible datacenter. As such, the flexible datacenter 500 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. And in still other embodiments, the flexible datacenter 500 might not include a mobile container. For example, the flexible datacenter 500 may be situated within a building or another type of stationary environment.

FIG. 6B shows the computing systems 512 in a straight-line configuration for installation within the flexible datacenter 500, according to one or more example embodiments. As indicated above, the flexible datacenter 500 may include a plurality of racks 604, each of which may include one or more computing systems 512 disposed therein. As discussed above, the power input system 502 may provide three phases of AC voltage to the power distribution system 506. In some examples, the power distribution system 506 may controllably provide a single phase of AC voltage to each computing system 512 or group of computing systems 512 disposed within the flexible datacenter 500. As shown in FIG. 6B, for purposes of illustration only, eighteen total racks 604 are divided into a first group of six racks 606, a second group of six racks 608, and a third group of six racks 610, where each rack contains eighteen computing systems 512. The power distribution system (506 of FIG. 5) may, for example, provide a first phase of three-phase AC voltage to the first group of six racks 606, a second phase of three-phase AC voltage to the second group of six racks 608, and a third phase of three-phase AC voltage to the third group of six racks 610. In other embodiments, the quantity of racks and computing systems can vary.

FIG. 7 shows a control distribution system 700 of the flexible datacenter 500 according to one or more example embodiments. The system 700 includes a grid operator 702, a generation station control system 216, a remote master control system 300, and a flexible datacenter 500. As such, the system 700 represents one example configuration for controlling operations of the flexible datacenter 500, but other configurations may include more or fewer components in other arrangements.

The datacenter control system 504 may independently or cooperatively with one or more of the generation station control system 414, the remote master control system 300, and/or the grid operator 702 modulate power at the flexible datacenter 500. During operations, the power delivery to the flexible datacenter 500 may be dynamically adjusted based on conditions or operational directives. The conditions may correspond to economic conditions (e.g., cost for power, aspects of computational operations to be performed), power-related conditions (e.g., availability of the power, the sources offering power), demand response, and/or weather-related conditions, among others.

The generation station control system 414 may be one or more computing systems configured to control various aspects of a generation station (not independently illustrated, e.g., 216 or 400). As such, the generation station control system 414 may communicate with the remote master

36

control system 300 over a networked connection 706 and with the datacenter control system 704 over a networked or other data connection 708.

As discussed with respect to FIGS. 2 and 3, the remote master control system 300 can be one or more computing systems located offsite, but connected via a network connection 710 to the datacenter control system 504. The remote master control system 300 may provide supervisory controls or override control of the flexible datacenter 500 or a fleet of flexible datacenters (not shown).

The grid operator 702 may be one or more computing systems that are configured to control various aspects of the power grid (not independently illustrated) that receives power from the generation station. The grid operator 702 may communicate with the generation station control system 300 over a networked or other data connection 712.

The datacenter control system 504 may monitor BTM power conditions at the generation station and determine when a datacenter ramp-up condition is met. The BTM power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost to obtain power), low priced power, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs. For example, a datacenter control system may analyze future workload and near term weather conditions at the flexible datacenter.

In some instances, the datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702 to go offline or reduce power. As such, the datacenter control system 504 may enable the power input system 502 to provide power to the power distribution system 506 to power the computing systems 512 or a subset thereof.

The datacenter control system 504 may optionally direct one or more computing systems 512 to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems 512 are configured to perform blockchain hashing operations, the datacenter control system 504 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems 512 may be configured to perform high-throughput computing operations and/or high performance computing operations.

The remote master control system 300 may specify to the datacenter control system 504 what sufficient behind-the-meter power availability constitutes, or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 500. In such circumstances, the datacenter control system 504 may provide power to only a subset of computing systems, or operate the plurality of computing systems in a

US 10,608,433 B1

37

lower power mode, that is within the sufficient, but less than full, range of power that is available. In addition, the computing systems 512 may adjust operational frequency, such as performing more or less processes during a given duration. The computing systems 512 may also adjust internal clocks via over-clocking or under-clocking when performing operations.

While the flexible datacenter 500 is online and operational, a datacenter ramp-down condition may be met when there is insufficient or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702. The datacenter control system 504 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by the remote master control system 300 or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently.

An operational directive may be based on current dispatch-ability, forward looking forecasts for when behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the generation station control system 414, the remote master control system 300, or the grid operator 702. For example, the generation station control system 414, the remote master control system 300, or the grid operator 702 may issue an operational directive to flexible datacenter 500 to go offline and power down. When the datacenter ramp-down condition is met, the datacenter control system 504 may disable power delivery to the plurality of computing systems (e.g., 512). The datacenter control system 504 may disable 714 the power input system 502 from providing power (e.g., three-phase nominal AC voltage) to the power distribution system 506 to power down the computing systems 512 while the datacenter control system 504 remains powered and is capable of returning service to operating mode at the flexible datacenter 500 when behind-the-meter power becomes available again.

While the flexible datacenter 500 is online and operational, changed conditions or an operational directive may cause the datacenter control system 504 to modulate power consumption by the flexible datacenter 500. The datacenter control system 504 may determine, or the generation station control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power the flexible datacenter 500. In such situations, the datacenter control system 504 may take steps to reduce or stop power consumption by the flexible datacenter 500 (other than that required to maintain operation of datacenter control system 504).

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically reduce or withdraw power delivery to one or more computing systems 512 to meet the dictate. The datacenter control system 504 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (e.g., 512) to reduce power consumption. The datacenter control system 504 may dynamically reduce the power consumption of one or more computing

38

systems by reducing their operating frequency or forcing them into a lower power mode through a network directive.

Similarly, the flexible datacenter 500 may ramp up power consumption based on various conditions. For instance, the datacenter control system 504 may determine, or the generation control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in greater power generation, availability, or economic feasibility. In such situations, the datacenter control system 504 may take steps to increase power consumption by the flexible datacenter 500.

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to increase power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically increase power delivery to one or more computing systems 512 (or operations at the computing systems 512) to meet the dictate. For instance, one or more computing systems 512 may transition into a higher power mode, which may involve increasing power consumption and/or operation frequency.

One of ordinary skill in the art will recognize that datacenter control system 504 may be configured to have a number of different configurations, such as a number or type or kind of the computing systems 512 that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available behind-the-meter power. As such, the datacenter control system 504 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 8 shows a control distribution system 800 of a fleet of flexible datacenters according to one or more example embodiments. The control distribution system 800 of the flexible datacenter 500 shown and described with respect to FIG. 7 may be extended to a fleet of flexible datacenters as illustrated in FIG. 8. For example, a first generation station (not independently illustrated), such as a wind farm, may include a first plurality of flexible datacenters 802, which may be collocated or distributed across the generation station. A second generation station (not independently illustrated), such as another wind farm or a solar farm, may include a second plurality of flexible datacenters 804, which may be collocated or distributed across the generation station. One of ordinary skill in the art will recognize that the number of flexible datacenters deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more example embodiments.

The remote master control system 300 may provide directive to datacenter control systems of the fleet of flexible datacenters in a similar manner to that shown and described with respect to FIG. 7, with the added flexibility to make high level decisions with respect to fleet that may be counterintuitive to a given station. The remote master control system 300 may make decisions regarding the issuance of operational directives to a given generation station based on, for example, the status of each generation station where flexible datacenters are deployed, the workload distributed across fleet, and the expected computational demand required for one or both of the expected workload and predicted power availability. In addition, the remote master control system 300 may shift workloads from the first plurality of flexible datacenters 802 to the second plurality of flexible datacenters 804 for any reason, including, for

example, a loss of BTM power availability at one generation station and the availability of BTM power at another generation station. As such, the remote master control system 300 may communicate with the generation station control systems 806A, 806B to obtain information that can be used to organize and distribute computational operations to the fleets of flexible datacenters 802, 804.

FIG. 9 shows a queue distribution arrangement for a traditional datacenter 902 and a flexible datacenter 500, according to one or more example embodiments. The arrangement of FIG. 9 includes a flexible datacenter 500, a traditional datacenter 902, a queue system 312, a set of communication links 916, 918, 920A, 920B, and the remote master control system 300. The arrangement of FIG. 9 represents an example configuration scheme that can be used to distribute computing operations using a queue system 312 between the traditional datacenter 902 and one or more flexible datacenters. In other examples, the arrangement of FIG. 9 may include more or fewer components in other potential configurations. For instance, the arrangement of FIG. 9 may not include the queue system 312 or may include routes that bypass the queue system 312.

The arrangement of FIG. 9 may enable computational operations requested to be performed by entities (e.g., companies). As such, the arrangement of FIG. 9 may use the queue system 312 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 500 and the critical traditional datacenter 902. Particularly, the arrangement of FIG. 9 may use the queue system 312 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations among datacenters. As a result, the use of the queue system 312 may reduce time to complete operations and reduce costs.

In some examples, one or more components, such as the datacenter control system 504, the remote master control system 300, the queue system 312, or the control system 936, may be configured to identify situations that may arise where using the flexible datacenter 500 can reduce costs or increase productivity of the system, as compared to using the traditional datacenter 902 for computational operations. For example, a component within the arrangement of FIG. 9 may identify when using behind-the-meter power to power the computing systems 512 within the flexible datacenter 500 is at a lower cost compared to using the computing systems 934 within the traditional datacenter 902 that are powered by grid power. Additionally, a component in the arrangement of FIG. 9 may be configured to determine situations when offloading computational operations from the traditional datacenter 902 indirectly (i.e., via the queue system 312) or directly (i.e., bypassing the queue system 312) to the flexible datacenter 500 can increase the performance allotted to the computational operations requested by an entity (e.g., reduce the time required to complete time-sensitive computational operations).

In some examples, the datacenter control system 504 may monitor activity of the computing systems 512 within the flexible datacenter 500 and use the respective activity levels to determine when to obtain computational operations from the queue system 312. For instance, the datacenter control system 504 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 512 to perform. The various factors may include power availability at the flexible datacenter 500 (e.g., either stored or from a BTM source), availability of the computing systems 512 (e.g., percentage of computing systems avail-

able), type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter 500, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system, among others. The datacenter control system 504 may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems 512 to perform. In such a configuration, the datacenter control system 504 manages the activity of the flexible datacenter 500, including determining when to acquire new sets of computational operations when capacity among the computing systems 512 permit.

In other examples, a component (e.g., the remote master control system 300) within the system may assign or distribute one or more sets of computational operations organized by the queue system 312 to the flexible datacenter 500. For example, the remote master control system 300 may manage the queue system 312, including the distribution of computational operations organized by the queue system 312 to the flexible datacenter 500 and the traditional datacenter 902. The remote master control system 300 may utilize to information described with respect to the Figures above to determine when to assign computational operations to the flexible datacenter 500.

The traditional datacenter 902 may include a power input system 930, a power distribution system 932, a datacenter control system 936, and a set of computing systems 934. The power input system 930 may be configured to receive power from a power grid and distribute the power to the computing systems 934 via the power distribution system 932. The datacenter control system 936 may monitor activity of the computing systems 934 and obtain computational operations to perform from the queue system 312. The datacenter control system 936 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 934 to perform. A component (e.g., the remote master control system 300) within the arrangement of FIG. 9 may assign or distribute one or more sets of computational operations organized by the queue system 312 to the traditional datacenter 902.

The communication link 916 represents one or more links that may serve to connect the flexible datacenter 500, the traditional datacenter 902, and other components within the system (e.g., the remote master control system 300, the queue system 312—connections not shown). In particular, the communication link 916 may enable direct or indirect communication between the flexible datacenter 500 and the traditional datacenter 902. The type of communication link 916 may depend on the locations of the flexible datacenter 500 and the traditional datacenter 902. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

The queue system 312 represents an abstract data type capable of organizing computational operation requests received from entities. As each request for computational operations are received, the queue system 312 may organize the request in some manner for subsequent distribution to a datacenter. Different types of queues can make up the queue system 312 within embodiments. The queue system 312 may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

In other examples, the queue system 312 may be distributed consisting of multiple queue sub-systems. In the distributed configuration, the queue system 312 may use multiple queue sub-systems to organize different sets of computational operations. Each queue sub-system may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required for performing each set of computational operations. For instance, a first queue sub-system may organize sets of non-intensive computational operations and a second queue sub-system may organize sets of intensive computational operations. In some examples, the queue system 312 may include queue sub-systems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations. In some examples, the queue system 312 may move computational operations between different computing systems or different datacenters in real-time.

Within the arrangement of FIG. 9, the queue system 312 is shown connected to the remote master control system 300 via the communication link 918. In addition, the queue system 312 is also shown connected to the flexible datacenter via the communication 920A and to the traditional datacenter 902 via the communication link 920B. The communication links 918, 920A, 920B may be similar to the communication link 916 and can be various types of communication links within examples.

The queue system 312 may include a computing system configured to organize and maintain queues within the queue system 312. In another example, one or more other components of the system may maintain and support queues within the queue system 312. For instance, the remote master control system 300 may maintain and support the queue system 312. In other examples, multiple components may maintain and support the queue system 312 in a distributed manner, such as a blockchain configuration.

In some embodiments, the remote master control system 300 may serve as an intermediary that facilitates all communication between flexible datacenter 500 and the traditional datacenter 902. Particularly, the traditional datacenter 902 or the flexible datacenter 500 might need to transmit communications to the remote master control system 300 in order to communicate with the other datacenter. As also shown, the remote master control system 300 may connect to the queue system 312 via the communication link 918. Computational operations may be distributed between the queue system 312 and the remote master control system 300 via the communication link 918. The computational operations may be transferred in real-time and mid-performance from one datacenter to another (e.g., from the traditional datacenter 902 to the flexible datacenter 500). In addition, the remote master control system 300 may manage the queue system 312, including providing resources to support queues within the queue system 312.

As a result, the remote master control system 300 may offload some or all of the computational operations assigned to the traditional datacenter 902 to the flexible datacenter 500. This way, the flexible datacenter 500 can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist traditional datacenter 902. The remote master control system 300 may use the queue system 312 to temporarily store and organize the

offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter 500) is available to perform them. The flexible datacenter 500 consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to the traditional datacenter 902. The remote master control system 300 may further communicate with the flexible datacenter 500 via communication link 922 and the traditional datacenter 902 via the communication link 924.

FIG. 10A shows method 1000 of dynamic power consumption at a flexible datacenter using behind-the-meter power according to one or more example embodiments. Other example methods may be used to manipulate the power delivery to one or more flexible datacenters.

In step 1010, the datacenter control system, the remote master control system, or another computing system may monitor behind-the-meter power availability. In some embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1020, the datacenter control system or the remote master control system 300 may determine when a datacenter ramp-up condition is met. In some embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the generation station to go offline or reduce power.

In step 1030, the datacenter control system may enable behind-the-meter power delivery to one or more computing systems. In some instances, the remote master control system may directly enable BTM power delivery to computing systems within the flexible system without instructing the datacenter control system.

In step 1040, once ramped-up, the datacenter control system or the remote master control system may direct one or more computing systems to perform predetermined computational operations. In some embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

While operational, the datacenter control system, the remote master control system, or another computing system may receive an operational directive to modulate power consumption. In some embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system or the remote master control system may dynamically reduce power delivery to one or more computing systems or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system or the remote master control system may dynamically adjust power delivery to one or more computing systems to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system may disable power delivery to one or more computing systems.

FIG. 10B shows method 1050 of dynamic power delivery to a flexible datacenter using behind-the-meter power according to one or more embodiments. In step 1060, the datacenter control system or the remote master control system may monitor behind-the-meter power availability. In

US 10,608,433 B1

43

certain embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1070, the datacenter control system or the remote master control system may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the generation station to go offline or reduce power.

In step 1080, the datacenter control system may disable behind-the-meter power delivery to one or more computing systems. In step 1090, once ramped-down, the datacenter control system remains powered and in communication with the remote master control system so that it may dynamically power the flexible datacenter when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system may dynamically modulate power delivery to one or more computing systems of a flexible datacenter based on behind-the-meter power availability or an operational directive. The flexible datacenter may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter may have a blackout state, where all power consumption, including that of the datacenter control system is halted. However, once the flexible datacenter enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system. Generation station conditions or operational directives may cause flexible datacenter to ramp-up, reduce power consumption, change power factor, or ramp-down.

FIG. 11 illustrates a block diagram of a system for implementing control strategies based on a power option agreement, according to one or more embodiments. The system 1100 represents an example arrangement that includes a control system (e.g., the remote master control system 262), a load (e.g., one or more of the datacenters 1102, 1104, and 1106), and a power entity 1140, which may establish and operate in accordance with a power option agreement. Additional arrangements are possible within examples.

In general, a power option agreement is an agreement between a power entity 1140 associated with the delivery of power to a load (e.g., a grid operator, power generation station, or local control station) and the load (e.g., the datacenters 1102-1106). As part of the power option agreement, the load (e.g., load operator, contracting agent for the load, semi-automated control system associated with the load, and/or automated control system associated with the load) provides the power entity 1140 with the right, but not obligation, to reduce the amount of power delivered (e.g., grid power) to the load up to an agreed amount of power during an agreed upon time interval. In order to provide the power entity 1140 with this option, the load needs to be using at least the amount of power subject to the option (e.g., a minimum power threshold). For instance, the load may agree to use at least 1 MW of grid power at all times during a specified 24-hour time interval to provide the power entity 1140 with the option of being able to reduce the amount of power delivered to the load by any amount up to 1 MW at any point during the specified 24-hour time interval. The load may grant the power entity 1140 with this option in exchange for a monetary consideration (e.g., receive power

44

at a reduced price and/or monetary payment if the option is exercised by the power entity).

The power option agreement may be used by the power entity 1140 to reserve the right to reduce the amount of grid power delivered to the load during a set time frame (e.g., the next 24 hours). For instance, the power entity 1140 may exercise a predefined power option to reduce the amount of grid power delivered to the load during a time when the grid power may be better redirected to other loads coupled to the power grid. As such, the power entity 1140 may exercise power option agreements to balance loads coupled to the power grid. In some embodiments, a power option agreement may also specify other parameters, such as costs associated with different levels of power consumption and/or maximum power thresholds for the load to operate according to.

To illustrate an example, a power option agreement may specify that a load (e.g., the datacenters 1102-1106) is required to use at least 10 MW or more at all times during the next 12 hours. Thus, the minimum power threshold according to the power option agreement is 10 MW and this minimum power threshold extends across the time interval of the next 12 hours. In order to comply with the agreement, the load must subsequently operate using 10 MW or more power at all times during the next 12 hours. This way, the load can accommodate a situation where the power entity 1140 exercises the option. Particularly, exercising the option may trigger the load to reduce the amount of power it consumes by an amount up to 10 MW at any point during the 12 hour interval. By establishing this power option agreement, the power entity 1140 can manipulate the amount of power consumed at the load during the next 12 hours by up to 10 MW if power needs to be redirected to another load or a reduction in power consumption is needed for other reasons.

In the example arrangement of the system 1100 shown in FIG. 11, one or more of the datacenters (e.g., the flexible datacenters 1102, 1104, and the traditional datacenter 1106) may operate as the load that is subject to a power option agreement. As the load that is subject to the power option agreement, the datacenters 1102-1106 may execute control instructions in accordance with power target consumption targets that meet or exceed the minimum power thresholds based on the power option agreement.

As shown in FIG. 11, each datacenter 1102-1106 may include a set of computing systems configured to perform computational operations using power from one or more power sources (e.g., BTM power, grid power, and/or grid power subject to a power option agreement). In particular, the flexible datacenter 1102 includes computing systems 1108 arranged into a first set 1114A, a second set 1114B, and a third set 1114C, the flexible datacenter 1104 includes computing systems 1110 arranged into a first set 1116A, a second set 1116B, and a third set 1118B, and the traditional datacenter 1106 includes computing systems 1112 arranged into a first set 1118A, a second set 1118B, and a third set 1118C. Each set of computing systems may include various types of computing systems that can operate in one or more modes.

The different sets of computing systems as well as the multiple datacenters are included in FIG. 11 for illustration purposes. In particular, the variety of computing systems represent different configurations that a load may take while operating in accordance with a power option agreement, and each configuration (as detailed herein) may include ramping up or down power consumption and transferring and performing computational operations between sets of comput-

US 10,608,433 B1

45

ing systems and/or datacenters. In other examples, the load that is subject to a power option agreement may take on other configurations (e.g., a single datacenter **1102-1106**, and/or a single set of computing systems).

The remote master control system **262** may serve as a control system that can determine performance strategies and provide control instructions to the load (e.g., one or more of the datacenters **1102-1106**). In particular, the remote master control system **262** can monitor conditions in concert with the minimum power thresholds and time intervals (e.g., power option data) set forth in, and/or derived from, one or more power option agreements to determine performance strategies that can enable the load to meet the expectations of the power option agreement(s) while also efficiently using power to accomplish computational operations. In some instances, the remote master control system **262** may also be subject to the power option agreement and may adjust its own power consumption based on the power option agreement (e.g., ramp up or down power consumption based on the defined minimum power thresholds during time intervals).

To establish a power option agreement, the remote master control system **262** (or another computing system) may communicate with the power entity **1140**. For instance, the remote master control system **262** may provide a request (e.g., a signal and/or a bid) to the power entity **1140** and receive the terms of one or more power option agreements, or power option data related to power option agreements (e.g., data such as minimum power thresholds and time intervals, but not all terms contained within a potential power option agreement) in response. In some examples, the remote master control system **262** may evaluate one or more conditions prior to establishing a power option agreement to ensure that the conditions could enable the load (e.g., the datacenters **1102-1106**) to operate in accordance with the power option agreement. For instance, the remote master control system **262** may check the quantity and deadlines associated with computational operations assigned to specific datacenters prior to establishing specific datacenters as a load subject to a power option agreement. In some cases, multiple power option agreements may be established. For example, each datacenter **1102-1106** may be subject to a different power option agreement, which may result in the remote master control system **262** managing the power consumption at each of the datacenters **1102-1106** differently.

Within the system **1100** shown in FIG. **11**, the power entity **1140** may represent any type of power entity associated with the delivery of power to the load that is subject to a power option agreement. For instance, the power entity **1140** may be a local station control system, a grid operator, or a power generation source. As such, the power entity **1140** may establish power option agreements with the loads via communication with the loads and/or the remote master control system **262**. For example, the power entity **1140** may obtain and accept a bid from a load trying to engage in a power option agreement with the power entity **1140**. The power entity **1140** is shown with a power option module **1142**, which may be used to establish power option agreements (e.g., fixed-duration **1144** and/or dynamic **1146**).

Once a power option agreement is established, the remote master control system **262** may obtain power option data from the power entity **1140** (or another source) that specifies the power and time expectations of the power entity **1140**. As shown in FIG. **11**, the power entity **1140** includes a power option module **1142**, which may be used to provide power option data to the remote master control system **262** and/or

46

the datacenters **1102-1106**. In particular, the power option data may specify the minimum power threshold or thresholds associated with one or more time intervals for the load to operate at in accordance with based on the power option agreement. The power option data may also specify other constraints that the load should operate in accordance with.

In some examples, the power option data may also include an indication of a monetary penalty that would be imposed upon the load for failure to operate as agreed upon for the power option agreement. In addition, the power option data may also include an indication of a monetary benefit provided to the load operating at power consumption levels that are in accordance with a power option agreement. For instance, monetary benefits could include reduced prices for power, credits for power, and/or monetary payments. In addition, the power option data may include further constraints upon power use, such as one or more maximum power thresholds and corresponding time intervals for the maximum power thresholds.

In some embodiments, the power entity **1140** may correspond to a qualified scheduling entity (QSE). A QSE may submit bids and offers on behalf of resource entities (REs) or load serving entities (LSEs), such as retail electric providers (REPs). QSEs may submit offers to sell and/or bids to buy power (energy) in the Day-Ahead Market (e.g., the next 24 hours) and the Real-Time Market. As such, the remote master control system **262** or another computing system may communicate with one or more QSEs to engage and control one or more loads in accordance with one or more power option agreements.

In some examples, a power option agreement may take the form of a fixed duration power option agreement **1144**. The fixed duration power option agreement **1144** may specify a set of minimum power thresholds and a set of time intervals in advance for an upcoming fixed duration of time covered by the agreement. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. Examples of such association are provided in FIG. **12**. The fixed duration power option agreement may be established in advance of the time period covered by the set of time intervals to enable the remote master control system **262** to prepare performance strategies for the load (e.g., the datacenter(s)) associated with the power option agreement. Thus, the remote master control system **262** may evaluate the fixed duration power option and other monitored conditions to determine performance strategies for a set of computing systems (e.g., one or more datacenters) during the different intervals that satisfy the minimum power thresholds.

In other examples, a power option agreement may take the form of a dynamic power option agreement **1146**. For a dynamic power option agreement **1146**, minimum power thresholds may be provided to the remote master control system **262** in real-time (or near real-time). For instance, a dynamic power option agreement may specify that the power entity **1140** may provide adjustments to minimum power thresholds and corresponding time intervals in real-time to the remote master control system **262**. For example, a dynamic power option agreement may provide power option data that specifies a minimum power threshold for immediate adjustments (e.g., for the next hour).

In an embodiment, a dynamic power option agreement **1146** may involve repeat communication between the remote master control system **262** and the power entity **1140**. Particularly, the power entity **1140** may provide signals to the remote master control system **262** that request power consumption adjustments to be initiated at one or more

datacenters by the remote master control system 262 over short time intervals, such as across minutes or seconds. For example, the power entity 1140 may communicate to the remote master control system 262 to ramp power consumption down to a particular level within the next 5 minutes. As a result, the remote master control system 262 may provide instructions to one or more datacenters to ramp down power consumption using a linear ramp over the next 5 minutes to meet the particular level specified by the power entity 1140. The remote master control system 262 may monitor the linear ramp down of power consumption and increase or decrease the rate that the datacenter(s) ramp down power use based on projections and updates received from the power entity 1140. As a result, although the ramp down of power consumption may initially be performed in a linear manner to meet a power target threshold, the remote master control system 262 may adjust the rate of power consumption decrease based on updates from the power entity 1140. For example, 25 percent of the overall power consumption ramp down may occur during a first period (e.g., 4 minutes 30 seconds) of the 5 minutes and the remaining 75 percent of the overall power consumption ramp down may occur during the remaining period of the 5 minutes (e.g., the final 30 seconds). The example percentages are included for illustration purposes and can vary within examples based on various parameters, such as additional communication (e.g., adjustments) provided by the power entity 1140.

In further examples, a power option agreement may operate similarly to both a fixed-duration 1144 and a dynamic power option agreement 1146. Particularly, power option data specifying minimum power thresholds and corresponding time intervals may be provided in advance for the entire fixed-duration of time (e.g., the next 24 hours). Additional power option data may then be subsequently provided enabling the remote master control system 262 to make one or more adjustments to accommodate any changes specified within the additional power option data. For instance, additional power option data may indicate that a power entity exercised its option to deliver less power to the load. As a result, the remote master control system may instruct the load to adjust power consumption based on the power entity reducing the power threshold minimum via exercising the option.

As indicated above, the remote master control system 262 may monitor conditions in addition to the constraints set forth in power option data received from the power entity 1140. Particularly, the remote master control system 262 may monitor and analyze a set of conditions (including the power option data) to determine strategies for assigning, transferring, and otherwise managing computational operations using the one or more datacenters 1102-1106. The determined strategies may enable efficient operation by the datacenters while also ensuring that the datacenters operate at target power consumption levels that meet or exceed the minimum power thresholds set forth within one or more power option agreements.

Example monitored conditions include, but are not limited to, power availability 1120, power prices 1122, computing systems parameters 1124, cryptocurrency prices 1126, computational operation parameters 1128, and weather conditions 1129. Power availability 1120 may include determining power consumption ranges at a set of computing systems and/or at one or more datacenters. In addition, power availability 1120 may also involve determining the source or sources of power available at a datacenter. For instance, the remote master control system 262 may identify the types of power sources (e.g., BTM, grid

power, and/or a battery system) that a datacenter has available. Power prices 1122 may involve an analysis of the different costs associated with powering a set of computing systems. For instance, the remote master control system 262 may determine cost of power from the grid without a power option agreement relative to the cost power from the grid under the power option agreement. In addition, the remote master control system 262 may also compare the cost of grid power relative to the cost of BTM power when available at a datacenter. The power prices 1122 may also involve comparing the cost of using power at different datacenters to determine which datacenter may perform computational operations at a lower cost.

Monitoring computing system parameters 1124 may involve determining parameters related to the computing systems at one or more datacenters. For instance, the remote master control system 262 may monitor various parameters of the computing systems at a datacenter, such as the abilities and availability of various computing systems, the status of the queue used to store computational operations awaiting performance by the computing systems. The remote master control system 262 may determine types and operation modes of the computing systems, including which computing systems could operate in different modes (e.g., a higher power or a lower power mode) and/or at different hash rates and/or frequencies. The remote master control system 262 may also estimate when computing systems may complete current computational operations and/or how many computational operations are assigned to computing systems.

Monitoring cryptocurrency prices 1126 may involve monitoring the current price of one or more cryptocurrencies, the hash rate and/or estimated power consumption associated with mining each cryptocurrency, and other factors associated with the cryptocurrencies. The remote master control system 262 may use data related to monitoring cryptocurrency prices 1126 to determine whether using computing systems to mine a cryptocurrency generates more revenue than the cost of power required for performance of the mining operations.

The remote master control system 262 may monitor parameters related to computational operations (e.g., computational operation parameters 1128). For example, the remote master control system 262 may monitor parameters related to the computational operations requiring performance and currently being performed, such quantity of operations, estimated time to complete, cost to perform each computational operation, deadlines and priorities associated with each computational operation. In addition, the remote master control system 262 may analyze computational operations to determine if a particular type of computing system may perform the computational operation better than other types of computing systems.

Monitoring weather conditions 1129 may include monitoring for any potential power generation disruption due to emergencies or other events, and changes in temperatures or weather conditions at power generators or datacenters that could affect power generation. As such, the operations and environment analysis module (or another component) of the remote master control system 262 may be configured to monitor one or more conditions described above.

The performance strategy determined by the remote master control system 262 based on the monitored conditions and/or power option data can include control instructions for the load (e.g., the datacenters and/or one or more sets of computing systems). For instance, a performance strategy can specify operating parameters, such as operating frequen-

cies, power consumption targets, operating modes, power on/off and/or standby states, and other operation aspects for computing systems at a datacenter.

The performance strategy can also involve aspects related to the assignment, transfer, and performance of computational operations at the computing systems. For instance, the performance strategy may specify computational operations to be performed at the computing systems, an order for completing computational operations based on priorities associated with the computational operations, and an identification of which computing systems should perform which computational operations. In some instances, priorities may depend on revenue associated with completing each computational operation and deadlines for each computational operation.

The monitored conditions may enable efficient distribution and performance of computational operations among computing systems at one or more datacenters (e.g., datacenters 1102-1106) in ways that can reduce costs and/or time to perform computational operations, take advantage of availability and abilities of computing systems at the datacenters 1102-1106, and/or take advantage in changes in the cost for power at the datacenters 1102-1106. In addition, the monitored conditions may also involve consideration of the power option data to ensure that the computing systems consume enough power to meet minimum power thresholds set forth in one or more power option agreements.

The various monitored conditions described above as well as other potential conditions may change dynamically and with great frequency. Thus, to enable efficient distribution and performance of the computational operations at the datacenters, the remote master control system 262 may be configured to monitor changes in the various conditions to assist with the efficient management and operations of the computing systems at each datacenter. For instance, the remote master control system 262 may engage in wired or wireless communication 1130 with datacenter control systems (e.g., datacenter control system 504) at each datacenter as well as other sources (e.g., the power entity 1140) to monitor for changes in the conditions.

The remote master control system 262 may analyze the different conditions in real-time to modulate operating attributes of computing systems at one or more of the datacenters. By using the monitored conditions, the remote master control system 262 may increase revenue, decrease costs, and/or increase performance of computational operations via various modifications, such as transferring computational operations between datacenters or sets of computing systems within a datacenter and adjusting performance at one or more sets of computing systems (e.g., switching to a low power mode).

In some examples, the traditional datacenter 1106 may be the load subject to a power option agreement. As such, the remote master control system 262 may factor the power option agreement when determining whether to perform computational operations using the computing systems 1112 at the traditional datacenter 1106 and/or transfer computational operations to the computing systems 1108, 1110 at the flexible datacenters 1102, 1104. For instance, the monitored conditions may indicate that the price of grid power is substantially higher than BTM power. As a result, the remote master control system 262 may transfer a subset of computational operations from the traditional datacenter 1106 to the flexible datacenters 1102, 1104. The traditional datacenter 1106 may still have some computational operations to perform to ensure that the traditional datacenter 1106 is

using enough power to meet the minimum power threshold or thresholds set forth in the power option agreement.

In some examples, the remote master control system 262 may monitor the grid frequency signal received from the power entity 1140. When the frequency of the grid deviates a threshold amount (e.g., 0.036 Hz above or below 60 Hz), the remote master control system 262 may adjust performance strategies at the load. In some cases, the remote master control system 262 may adjust the power consumption at the load, the number of miners (or computing systems) operating at the load, and/or the frequency or hash rate, among other possible changes. The remote master control system may readjust performance strategies at the load in response to receiving additional power option data from the power entity 1140 (e.g., an indication that the frequency of the grid is back to 60 Hz). In addition, the remote master control system 262 may communicate changes in operations at the load to the power entity 1140. This way, the power entity 1140 may obtain confirmation that the load is adjusting in accordance with a power option agreement.

In some embodiments, a power generation source (e.g., the generation station 400 shown in FIG. 4) may enter into a power option agreement with a grid operator, which may provide the grid operator with the option to reduce the amount of power that the power source generator can deliver to the grid during a defined time interval. For instance, a wind generation farm may enter into the power option agreement with the grid operator. In addition, the remote master control system 262 may also enter into a power option agreement with the power generation source (e.g., the wind farm) to provide a load that can receive excess power from the power generation source when the grid operator exercises the option and lowers the amount of power that the power generation source can deliver to the grid. Thus, rather than reducing the amount of power produced, the power generation source could exercise an option in the agreement with remote master control system 262 and redirect excess power to one or more loads (e.g., a set of computing systems) that could ramp up power consumption in response. In such situations, the remote master control system 262 maybe able to use the excess power from the power generation source (e.g., BTM power) to perform operations at one or more loads at a low cost (or no cost at all). In addition, the power generation source may benefit from the power option agreement by directing excess power to the load instead of temporarily halting power production.

In some examples, a power option agreement may depend on parameters associated balancing grid capacity and demand. For instance, power option agreements may incentivize power consumption ramping during periods of peak grid power use.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments. The graph 1200 shows power option data arranged according to power 1204 over time 1202. As shown in FIG. 12, time 1202 increases along the X-axis and minimum power thresholds 1204 increase along the Y-axis of the graph 1200. In the example embodiment shown in FIG. 12, the time 1202 increases up to a full day (e.g., 24 hours) in 4 hour increments and the power is shown in MW increasing in intervals of 5 MW. The 24 duration and example minimum power thresholds can differ in other embodiments. Particularly, these values may depend on the terms set forth within the power option agreement.

The graph line 1206 represents sets of minimum power thresholds 1206A, 1206B, 1206C that are specified by

US 10,608,433 B1

51

power option data based on the power option agreement. As shown, the graph line **1206** extends the entire 24 hour duration, which indicates that the set of time intervals associated with minimum power thresholds add up to 24 hours. In other examples, the power option agreement may not include a minimum power threshold during a portion of the duration.

The graph line **1206** of the graph **1200** is further used to illustrate power consumption levels that one or more loads (e.g., a set of computing systems) operating according to the power option agreement may utilize during the 24 hour duration. Particularly, the power quantities above the graph line **1206** represents power levels that the load(s) may consume from the power grid during the 24 hour duration that would satisfy the requirements (i.e., the minimum power thresholds **1206A-1206C**) set forth by the power option agreement. In particular, the power quantities above the graph line **1206** include any power quantity that meets or exceeds the minimum power threshold at that time. By extension, the power quantities positioned below the graph line **1206** represents the amount of power that the load could be directed to reduce power consumption by per the power option agreement.

To further illustrate, an initial minimum power threshold **1206A** is shown associated with the time interval starting at hour 0 and extending to hour 8. In particular, the minimum power threshold **1206A** is set at 5 MW during this time interval. Thus, based on the power option data shown in FIG. **12**, the loads must be able to operate at a target power consumption level that is equal to or greater than the 5 MW minimum power threshold **1206A** at all times during the time interval extending from hour 0 to hour 8, in order to be able to satisfy the power option if it is exercised for that time interval. Similarly, the power entity could reduce the power consumed by loads by any amount up to 5 MW at any point during the time interval from hour 0 to hour 8 in accordance with the power option agreement. For instance, the power entity could exercise its option at any point during this time interval to reduce the power consumed by the loads by 3 MW as a way to load balance the power grid. In response to the power entity exercising its option, the load may then operate using 3 MW less power and/or another strategy determined by a control system factoring additional conditions (e.g., the price of grid power, the revenue that could be generated from mining a cryptocurrency, and/or parameters associated with computational operations awaiting performance).

As further shown in the graph **1200** illustrated in FIG. **12**, the next minimum power threshold **1206B** is associated with the following time interval, which starts at hour 8 and extends until hour 16. During this time interval (hour 8 to hour 16), the load(s) may consume 10 MW or more power since the minimum power threshold **1206B** is now set at 10 MW as shown on the Y-axis of the graph **1200**. In light of the power option data, a control system may determine and provide a performance strategy to the load (e.g., a set of computing systems) that includes a power consumption target that meets or exceeds the minimum power threshold **1206B** (i.e., 10 MW). The performance strategy may depend on the power option data as well as other possible conditions, such as the price of grid power, the availability of computing systems, and/or the type of computing operations, etc. In addition, the power entity could exercise its option to reduce the amount of power consumed by the load by 10 MW or less as represented by the power levels under the minimum threshold **1206B** that extend during the time interval of hour 8 to hour 16.

52

The last minimum power threshold **1206C** is associated with the time interval that starts at hour 16 and extends until hour 24. Similar to the initial minimum power threshold **1206A** associated with the beginning of the graph line **1206**, the last minimum power threshold **1206** is also set at 5 MW. As such, at any point during this interval (hour 16 to hour 24) the loads may consume 5 MW or more to operate in accordance with the power option agreement. As discussed above, by operating at 5 MW or more, the load enables the power consumed from the power grid to be reduced any amount from zero up to 5 MW during this time interval.

When determining the power consumption strategy for a load, a computing system (e.g., the remote master control system **262**) may consider various conditions in addition to the power option data received based on one or more power option agreements. Particularly, the computing system may consider and weigh different conditions in addition to the power option data to determine power consumption targets and/or other control instructions for a load. The conditions may include, but are not limited to, the price of grid power, the price of alternative power sources (e.g., BTM power, stored energy), the revenue associated with mining for one or more cryptocurrencies, parameters related to the computational operations requiring performance (e.g., priorities, deadlines, status of the queue organizing the operations, and/or revenue associated with completing each computational operation), parameters related to the set of computing systems (e.g., types and availabilities of computing systems), and other conditions (e.g., penalties if a minimum power threshold is not met and/or monetary benefits from operating under a power option agreement). By weighing various conditions, the computing system may efficiently manage the set of computing systems, including enabling performance of computational operations cost effectively and/or ensuring that computing systems operate at target power consumption levels that one or more satisfy power option agreements.

In some examples, the computing system may decrease the amount of power that a set of computing systems consumes from one source and while also increasing the amount of power that the set consumes from another source. For instance, the computing system may determine that the price of power grid power is above a threshold price that makes computational operations relatively expensive to perform using grid power. As a result, the computing system may provide control instructions for the computing systems to consume power grid power that matches a minimum power threshold specified by power option data. This may enable the computing systems to satisfy the power option agreement while also avoiding using pricey grid power beyond the minimum amount required per the power option data. In addition, the computing system may instruct some computing systems to switch to a low power mode or temporarily stop until the price of power from the grid decreases. The computing system may instruct one or more computing systems to operate using power from another source (e.g., BTM power and/or stored energy from a battery system) and/or transfer one or more computational operations to another set of computing systems (e.g., a different datacenter).

When the power option agreement is a fixed duration power option agreement, the computing system may receive an indication of all the minimum power thresholds **1206A-1206C** and an indication of the associated time interval altogether and in advance of the duration associated with the power option agreement. By providing all of the minimum power thresholds **1206A-1206C** and the time intervals in

US 10,608,433 B1

53

advance, the computing system may determine a performance strategy for the load that can extend across the entire duration. Particularly, the computing system may factor the minimum power thresholds and associated time intervals as well as other monitored conditions to determine the performance strategy for the total duration. This can enable the computing system to accept and assign computational operations to computing systems in advance while also using a performance strategy that meets the expectations of a power option agreement.

In some examples, the performance strategy determined by the computing system may include control instructions for the set of computing systems to execute if a power option is exercised. For instance, the performance strategy may specify different power consumption targets for the computing systems that depend on whether a power option is exercised during each time interval.

In some instances, the computing system may modify the performance strategy when one or more conditions change enough to warrant a modification. For instance, the computing system may receive an indication of a change in a minimum power threshold (e.g., a decrease in the minimum power threshold) and determine one or more modifications based on the new minimum power threshold and/or other conditions (e.g., a change in the price of power).

In other examples, the power option agreement may be a dynamic power option agreement. Particularly, the load may be subject to a changing minimum power threshold that can vary during a predefined duration associated with the power option agreement. For example, a dynamic power option agreement may specify that the load is subject to a minimum power threshold that may vary from 0 MW up to 5 MW during the next 24 hours and the particular minimum threshold for each hour may depend on power option data received from the power entity during the prior hour. The dynamic power option agreement may further specify the expected response time from the load. For instance, the power option agreement may indicate that an indication of a new minimum power threshold will be provided an hour prior to the start of the minimum power threshold. The computing system, for example, may receive an indication at hour 7 about the increase in the minimum power threshold **1206B** starting at hour 8. The indication may (or may not) specify the total time interval associated with a new minimum power threshold. For instance, the indication received by the computing system may specify that the 10 MW minimum power threshold **1206B** extends from hour 8 until hour 16. In other instances, the power option data may indicate that the computing system should abide by the new minimum power threshold until receiving further power option data indicating a change to another new minimum power threshold.

In some examples, the power option data may arrive at the computing system in an unknown order from the power entity with expectations of swift power consumption adjustments by the load. As a result, the power option agreement may require fast ramping of the load to meet changes. Ramping may involve ramping up or down power consumption as well as ramping operating techniques (e.g., adjusting frequency or operation mode).

In some embodiments, the type of power option power agreement may depend on the delivery and content of power option data provided to the load (or a control system controlling the load). For instance, a computing system may receive minimum power thresholds set across an entire duration associated with a power option agreement in advance when the power option agreement is a fixed-

54

duration power option agreement. In other instances, the computing system may receive power option data dynamically and adjust operations in real-time (or near real-time). For instance, the computing system may receive a series of power option data that each specifies minimum power threshold changes during the duration set forth in the dynamic power option agreement. To illustrate an example, the computing system may receive power option data during hour 1 that specifies the minimum power threshold for hour 2, power option data during hour 2 that specifies the minimum power threshold for hour 3, and so on across the duration of the dynamic power option agreement.

In some examples, the minimum power threshold for a time interval may be zero during the duration of a power option agreement. As such, the load may use any amount of power from the power grid in accordance with the power option agreement, including no power at all during this time interval. When the price for power is high during this time frame, the load may ramp down power usage to zero MW to avoid paying the high price for power while still being in compliance with the power option agreement.

FIG. 13 illustrates a method for implementing control strategies based on a fixed-duration power option agreement, according to one or more embodiments. The method **1300** serves as an example and may include other steps within other embodiments. A control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1300**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At step **1302**, the method **1300** involves monitoring a set of conditions. For instance, a computing system (e.g., a control system) may monitor various conditions that could impact the performance of operations at one or more loads, including the power consumption targets at the loads. The set of monitored conditions may include a variety of information obtained from one or more external sources, such as one or more datacenters, databases, power generation stations, or types of sources.

Some example conditions include, but are not limited to, the price of grid power, the price and availability of alternative power options (e.g. BTM power, and/or stored energy), parameters of the load (e.g., ramping abilities, type of computing systems, operation modes, etc.), parameters of tasks to be performed using the power at the load (e.g., types, deadlines, priorities, and/or revenue associated with computational operations), availability of other computing systems and their associated costs, and/or revenue associated with mining a cryptocurrency. The computing system may monitor one or more of these conditions as well as others.

At step **1304**, the method **1300** involves receiving power option data based, at least in part, on a power option agreement. As discussed above, the computing system (e.g., a remote master control system) may engage in a power option agreement with a power entity. As a result, the computing system may control a load (e.g., a set of computing systems) in accordance with power thresholds and time intervals received from the power entity based on the power option agreement.

In some examples, the power option data may specify a set of minimum power thresholds and a set of time intervals. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. To illustrate an example, the power option data may specify a first minimum power threshold

US 10,608,433 B1

55

associated with a first time interval and a second minimum power threshold associated with a second time interval, with the second time interval subsequent to the first time interval.

The set of time intervals may add up to the duration represented by the power option agreement. For instance, the total duration of the set of time intervals may correspond to a twenty-four hour period (e.g., the next day). In other examples, the power option agreement may span across a different duration (e.g., 12 hours). In additional embodiments, the power option data may specify other information, such as monetary incentives associated with parameters of the power option agreement and/or one or more maximum power thresholds.

At step **1306**, the method **1300** involves determining a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy may be determined responsive to receiving the power option data. In addition, the performance strategy may include a power consumption target for the set of computing systems for each time interval in the set of time intervals. In some examples, each power consumption target is equal to or greater than the minimum power threshold associated with each time interval.

As an example, the performance strategy may specify a first power consumption target for the set of computing systems for a first time interval such that the first power consumption target is equal to or greater than a first minimum power threshold associated with the first time interval and a second power consumption target for the set for a second time interval in a similar manner (i.e., the second power consumption target is equal to or greater than a second minimum power threshold).

In some examples, the performance strategy may include an sequence for the set of computing systems to follow when performing computational operations. The sequence, for example, may be based on priorities associated with the computational operations. In addition, the performance strategy may include one or more power consumption targets that are greater than the minimum power thresholds when the price of power from the power grid is below a threshold price during the time intervals associated with the minimum power thresholds.

The performance strategy may also involve transferring, delaying, or adjusting one or more computational operations performed at the set of computing systems. In addition, the performance strategy may involve adjusting operations at the computing systems. For instance, one or more computing systems may switch modes (e.g., operate at a higher frequency or switch to a low power mode).

In addition, the performance strategy may also specify power consumption targets for the set of computing systems to use if the power option is exercised during an interval. This way, the computing systems may continue to perform computational operations (or suspend performance) based on the power option being exercised.

At step **1308**, the method **1300** involves providing instructions to the set of computing systems to perform one or more computational operations based on the performance strategy. For example, the set of computing systems may operate according to the performance strategy to ensure that the minimum power thresholds are met during the defined time intervals based on the power option agreement.

Some examples may further involve receiving subsequent power option data based, at least in part, on the power option agreement. The subsequent power option data may specify to decrease one or more minimum power thresholds of the

56

set of power thresholds. Responsive to receiving the subsequent power option data, the performance strategy for the set of computing systems may be modified based on a combination of at least a portion of the subsequent power option data and one or more conditions of the monitored conditions. The modified performance strategy may include one or more reduced power consumption targets for the set of computing systems. The amount of the reduction in a power consumption target may depend linearly with the amount that the corresponding minimum power threshold was reduced by. For instance, when a minimum power threshold for a time interval is reduced from 10 MW to 5 MW, the power consumption target for that time interval may be reduced from 10 MW to 5 MW. Instructions may be provided to the set of computing systems to perform computational operations based on the modified performance strategy.

FIG. 14 illustrates a method for implementing control strategies based on a dynamic power option agreement, according to one or more embodiments. The method **1400** serves as an example and may include other steps within other embodiments. Similar to the method **1400**, a control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1400**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At block **1402**, the method **1400** involves monitoring a set of conditions. Similar to block **1302** of the method **1300**, a computing system may monitor various conditions to determine instructions for controlling a set of computing systems.

At block **1404**, the method **1400** involves receiving first power option data based, at least in part, on a power option agreement while monitoring the set of conditions. The first power option data may specify a first minimum power threshold associated with a first time interval. For example, the first power option data may specify a minimum power threshold of 10 MW for the next hour, which may start in an hour or less.

The power option agreement may correspond to a dynamic power option agreement in some examples. When managing a load with respect to a dynamic power option agreement, a computing system may receive power option data specifying changes in minimum power thresholds that a load (e.g., the set of computing systems) may be designated to use in the near term (e.g., the next hour). For example, the computing system may receive power option data during each hour of the duration specified by a power option agreement that indicates a minimum power threshold for the next hour.

At block **1406**, the method **1400** involves providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition. The first control instructions may be provided responsive to receiving the first power option data.

The first control instructions may include a first power consumption target for the set of computing systems for the first time interval. Particularly, the first power consumption target may be equal to or greater than the first minimum power threshold associated with the first time interval. For example, the first power consumption target may be greater than the first minimum power threshold when a cost of power from the power grid is below a threshold price during the first time interval. In other instances, the first power consumption target may be equal to the first minimum power

US 10,608,433 B1

57

threshold when the cost of power from the power grid is greater than the threshold price.

In some examples, control instructions may specify a sequence for the computing systems to follow when performing computational operations. The sequence may be based on priorities associated with each computational operation.

The first control instructions may be determined based on a combination of the first power option data, the price of power from the power grid, and parameters associated with computational operations to be performed at the set of computing systems.

In some examples, the first control instructions may involve ramping up or down power consumption at the set of computing systems. The power consumption may be ramped up or down based on the first minimum power threshold and one or more other conditions (e.g., the price of power).

At block 1408, the method 1400 involves receiving second power option data based, at least in part, on the power option agreement while monitoring the set of conditions. The computing system may receive the second power option data subsequent to receiving the first power option data. The second power option data may specify a second minimum power threshold associated with a second time interval. For example, the second minimum power threshold may be 7 MW over the duration of the upcoming hour. In other examples, the second minimum power threshold may differ as shown in FIG. 12.

In some instances, the computing system may receive the second power option data during the first time interval such that the second time interval overlaps the first time interval. For instance, the computing system may receive the second power option data to enable real-time adjustments to be made to the power consumed at the set of computing systems.

At block 1410, the method 1400 involves providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power option data and at least one condition. The second control instructions may be provided responsive to receiving the second power option data. The second control instructions may specify a second power consumption target for the set of computing systems for the second time interval. The second power consumption target may be equal to or greater than the second minimum power threshold associated with the second time interval.

In some examples, the computing system may provide a request to a QSE to determine the power option agreement. As such, the computing system may receive power option data (e.g., the first and second power option data) in response to providing the request to the QSE.

The computing system may monitor the price of power from the power grid, and the global mining hash rate and a price for a cryptocurrency (e.g., Bitcoin), among other conditions. The computing system may determine control instructions (e.g., the first and/or second control instructions) based on a combination of power option data, the price of power from the power grid, and the global mining hash rate and the price for the cryptocurrency. For instance, the computing system may cause one or more computing systems (e.g., a subset of computing systems) to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

58

Advantages of one or more embodiments of the present invention may include one or more of the following:

One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive low cost or unutilized power behind-the-meter when it is available.

One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations across multiple datacenters.

One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

One or more embodiments of the present invention allows for management and distribution of computational operations at computing systems across a fleet of datacenters such that the performance of the computational operations take advantages of increased efficiency and decreased costs.

It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

US 10,608,433 B1

59

What is claimed is:

1. A system comprising:
 - a set of computing systems, wherein the set of computing systems is configured to perform computational operations using power from a power grid;
 - a control system configured to:
 - monitor a set of conditions;
 - receive power option data based, at least in part, on a power option agreement, wherein the power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, wherein each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals;
 - responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions, wherein the performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, wherein each power consumption target is equal to or greater than the minimum power threshold associated with each time interval; and
 - provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.
2. The system of claim 1, wherein the control system is configured to monitor the set of conditions comprising:
 - a price of power from the power grid; and
 - a plurality of parameters associated with one or more computational operations to be performed at the set of computing systems.
3. The system of claim 2, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at least the portion option data, the price of power from the power grid, and the plurality of parameters associated with the one or more computational operations.
4. The system of claim 3, wherein the performance strategy further comprises:
 - an order for the set of computing systems to follow when performing the one or more computational operations, wherein the order is based on respective priorities associated with the one or more computational operations.
5. The system of claim 4, wherein the performance strategy further comprises:
 - at least one power consumption target that is greater than a minimum power threshold when the price of power from the power grid is below a threshold price during the time interval associated with the minimum power threshold.
6. The system of claim 1, wherein the control system is further configured to:
 - receive subsequent power option data based, at least in part, on the power option agreement,
 - wherein the subsequent power option data specify to decrease one or more minimum power thresholds of the set of minimum power thresholds.
7. The system of claim 6, wherein the control system is further configured to:
 - responsive to receiving the subsequent power option data,
 - modify the performance strategy for the set of computing systems based on a combination of at least the

60

- portion of the subsequent power option data and at least one condition in the set of conditions,
- wherein the modified performance strategy comprises one or more reduced power consumption targets for the set of computing systems.
8. The system of claim 7, wherein the control system is further configured to:
 - provide instructions to the set of computing systems to perform the one or more computational operations based on the modified performance strategy.
9. The system of claim 1, wherein the control system is a remote master control system positioned remotely from the set of computing systems.
10. The system of claim 1, wherein the control system is a mobile computing device.
11. The system of claim 1, wherein the control system is configured to receive the power option data while monitoring the set of conditions.
12. The system of claim 1, wherein the control system is further configured to:
 - provide a request to a qualified scheduling entity (QSE) to determine the power option agreement; and
 - receive power option data in response to providing the request to the QSE.
13. The system of claim 1, wherein the power option data specify: (i) a first minimum power threshold associated with a first time interval in the set of time intervals, and (ii) a second minimum power threshold associated with a second time interval in the set of time intervals,
 - wherein the second time interval is subsequent to the first time interval.
14. The system of claim 13, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems such that the performance strategy comprises:
 - a first power consumption target for the set of computing systems for the first time interval, wherein the first power consumption target is equal to or greater than the first minimum power threshold; and
 - a second power consumption target for the set of computing systems for the second time interval, wherein the second power consumption target is equal to or greater than the second minimum power threshold.
15. The system of claim 1, wherein a total duration of the set of time intervals corresponds to a twenty-four hour period.
16. The system of claim 1, wherein the set of conditions monitored by the control system further comprise:
 - a price of power from the power grid; and
 - a global mining hash rate and a price for a cryptocurrency; and
 wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at the portion of the power option data, the price of power from the power grid, the global mining hash rate and the price for the cryptocurrency,
 - wherein the performance strategy specifies for at least a subset of the set of computing systems to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

61

17. A method comprising:
monitoring, by a computing system, a set of conditions;
receiving, at the computing system, power option data
based, at least in part, on a power option agreement,
wherein the power option data specify: (i) a set of
minimum power thresholds, and (ii) a set of time
intervals, wherein each minimum power threshold in
the set of minimum power thresholds is associated with
a time interval in the set of time intervals;
responsive to receiving the power option data, determin-
ing a performance strategy for a set of computing
systems based on a combination of at least a portion of
the power option data and at least one condition in the
set of conditions, wherein the performance strategy
comprises a power consumption target for the set of
computing systems for each time interval in the set of
time intervals, wherein each power consumption target
is equal to or greater than the minimum power thresh-
old associated with each time interval; and
providing instructions to the set of computing systems to
perform one or more computational operations based
on the performance strategy.
18. The method of claim 17, wherein determining the
performance strategy for the set of computing systems
comprises:
identifying information about the set of computing sys-
tems; and
determining the performance strategy to further comprise
instructions for at least a subset of the set of computing
systems to operate at an increased frequency based on
a combination of at least the portion of the power
option data and the information about the set of com-
puting systems.
19. The method of claim 17, further comprising:
receiving subsequent power option data based, at least in
part, on the power option agreement, wherein the
subsequent power option data specify to decrease one
or more minimum power thresholds of the set of
minimum power thresholds;

62

responsive to receiving the subsequent power option data,
modifying the performance strategy for the set of
computing systems based on a combination of at least
the portion of the subsequent power option data and at
least one condition in the set of conditions, wherein the
modified performance strategy comprises one or more
reduced power consumption targets for the set of com-
puting systems; and
providing instructions to the set of computing systems to
perform the one or more computational operations
based on the modified performance strategy.
20. A non-transitory computer readable medium having
stored therein instructions executable by one or more pro-
cessors to cause a computing system to perform functions
comprising:
monitoring a set of conditions;
receiving power option data based, at least in part, on a
power option agreement, wherein the power option
data specify: (i) a set of minimum power thresholds,
and (ii) a set of time intervals, wherein each minimum
power threshold in the set of minimum power thresh-
olds is associated with a time interval in the set of time
intervals;
responsive to receiving the power option data, determin-
ing a performance strategy for a set of computing
systems based on a combination of at least a portion of
the power option data and at least one condition in the
set of conditions, wherein the performance strategy
comprises a power consumption target for the set of
computing systems for each time interval in the set of
time intervals, wherein each power consumption target
is equal to or greater than the minimum power thresh-
old associated with each time interval; and
providing instructions to the set of computing systems to
perform one or more computational operations based
on the performance strategy.

* * * * *

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

**Title of
Invention**

Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration Power Option Agreement

As the below named inventor, I hereby declare that:

This declaration
is directed to:

☒

The attached application, or

☐

United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.

WARNING:

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Michael T. McNamara

Date (Optional): Dec 2, 2019

Signature: 

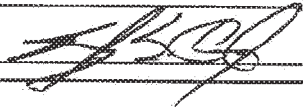
Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN APPLICATION DATA SHEET (37 CFR 1.76)

Title of Invention	Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration Power Option Agreement	
<p>As the below named inventor, I hereby declare that:</p> <p>This declaration is directed to: <input checked="" type="checkbox"/> The attached application, or <input type="checkbox"/> United States application or PCT international application number _____ filed on _____.</p> <p>The above-identified application was made or authorized to be made by me.</p> <p>I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.</p> <p>I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.</p> <p style="text-align: center;">WARNING:</p> <p>Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.</p>		
<p>LEGAL NAME OF INVENTOR</p> <p>Inventor: <u>Raymond E. Cline Jr.</u> Date (Optional): <u>12/3/2019</u></p> <p>Signature: </p>		
<p>Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.</p>		

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,

Plaintiffs,

V.

LANCIUM LLC, MICHAEL T.
MCNAMARA, and RAYMOND E. CLINE,
JR.

Defendants.

C.A. No. 21-534-MN

JURY TRIAL DEMANDED

AMENDED COMPLAINT

Plaintiffs BearBox LLC (“BearBox”) and Austin Storms (collectively, “Plaintiffs”) bring this action against Lancium LLC (“Lancium”), Michael T. McNamara, and Raymond E. Cline, Jr. (collectively “Defendants”) to correct the inventorship of U.S. Patent No. 10,608,433 (the “433 Patent”) and to recover damages, injunctive relief, declaratory relief, and other remedies for Defendants’ wrongful actions to obtain, misuse, disclose, and claim as their own Plaintiffs’ proprietary cryptocurrency mining technology. Plaintiffs further allege as follows:

INTRODUCTION

1. This case is about the Defendants' theft of inventions that rightfully belong to Plaintiffs.
2. ~~Plaintiffs developed~~Plaintiffs developed proprietary technology relating to cryptocurrency mining systems (the "BearBox Technology"). By way of background, the BearBox Technology generally relates to an energy-efficient cryptocurrency mining system and related methods that reduce the inefficiency and environmental impact of energy-expensive mining operations by better utilizing available energy resources to increase stability of the

energy grid, minimize a mining operation's impact on peak-demand, and also alleviate ~~electricity undersupply and/or oversupply conditions (the "BearBox Technology")~~. energy over-supply conditions. The BearBox Technology can be used to mine cryptocurrency, such as Bitcoin.

3. The Defendants induced the Plaintiffs to disclose the BearBox Technology to them under the guise of a possible business deal between Defendants and Plaintiffs to jointly commercialize the BearBox Technology. Before disclosing the BearBox Technology to Defendants, Plaintiffs obtained assurances of confidentiality from Defendants.

4. ~~Rather than keeping the BearBox Technology confidential and using it in furtherance of a business relationship with Plaintiffs, the~~ The Defendants stole the BearBox Technology from Plaintiffs by converting and misappropriating it and claiming it as their own. Defendants filed a U.S. patent application that wrongfully disclosed the BearBox Technology to the U.S. Patent and Trademark Office and ultimately to the public. The claimed subject matter of the '433 Patent falls fully within the scope of the BearBox Technology. And by obtaining the '433 Patent with claims directed to the BearBox Technology, the Defendants have wrongfully obtained a patent covering the BearBox Technology and wrongfully claimed the BearBox Technology as their own.

5. Plaintiffs bring this action to correct the named inventors on the '433 Patent. The inventions claimed in the '433 Patent are inventions conceived by Storms, founder and president of BearBox.

PARTIES

6. Plaintiff BearBox LLC ("BearBox") is a limited liability company organized and existing under the laws of Louisiana with its principal place of business at 4422 Highway 22, Mandeville, Louisiana 70471.

7. Plaintiff Austin Storms is an individual residing in Mandeville, Louisiana.

8. On information and belief, Defendant Lancium is a Delaware limited liability company with its principal place of business at 6006 Thomas Rd, Houston, Texas 77041. On information and belief, Lancium has a registered agent capable of accepting service in this district, Harvard Business Services, Inc. with a place of business at 16192 Coastal Highway, Lewes, DE 19958.

9. On information and belief, Defendant Michael T. McNamara is the Chief Executive Officer and a founder of Lancium and resides in Newport Beach, California. Defendant McNamara is named as a purported inventor on the face of the '433 Patent.

10. On information and belief, Defendant Raymond E. Cline, Jr. is the Chief Computing Officer of Lancium and resides in Houston, Texas. Defendant Cline is named as a purported inventor on the face of the '433 Patent.

JURISDICTION

11. This is an action seeking correction of the named inventors of a United States patent under 35 U.S.C. § 256. As such, this action arises under the laws of the United States.

12. This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under an Act of Congress relating to patents, specifically 35 U.S.C. § 256.

~~13. This Court further has subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under the trade secret laws of the United States, 18 U.S.C. § 1836.~~

∴

~~14.~~13. The Court has supplemental jurisdiction under 28 U.S.C. § 1367 over all asserted claims under state law because those claims are so related to the claims in this action that arise under federal law that they form part of the same case or controversy.

~~15.~~14. The Court also has jurisdiction pursuant to 28 U.S.C. § 1332, as complete diversity of citizenship exists among the parties, and the amount in controversy exceeds \$75,000. Plaintiff BearBox is a citizen of the State of Louisiana because it is organized under the laws of the State of Louisiana and has its principal place of business in the State of Louisiana. Plaintiff Storms is a citizen of the State of Louisiana because he resides in the State of Louisiana. In contrast, none of the Defendants are citizens of the State of Louisiana. Defendant Lancium is a citizen of the States of Delaware and Texas because it is organized under the laws of the State of Delaware and has its principal place of business in the State of Texas. Defendant McNamara is a citizen of the State of California because he resides in the State of California. Defendant Cline is a citizen of the State of Texas because he resides in the State of Texas. Therefore, because the Plaintiffs are both citizens of the State of Louisiana (and no other states) for purposes of diversity jurisdiction, and none of the Defendants are citizens of the State of Louisiana, complete diversity exists among the parties.

~~16.~~15. This Court has general personal jurisdiction over Lancium because it is organized under the laws of the State of Delaware and because it maintains an ongoing presence in this District at least through its registered agent.

~~17.~~16. This Court has specific personal jurisdiction over each of Defendants McNamara and Cline at least under Title 6 of the Delaware Code, § 18-109(a).

~~18.~~17. On information and belief, Defendant McNamara is the Chief Executive Officer of Lancium. On information and belief, as the Chief Executive Offer, McNamara participates

materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

~~19-18.~~ 19-18. McNamara is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from the interests of Lancium and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claims against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant McNamara together. Plaintiffs' claims against Defendant McNamara arise out of his exercise of his powers as Chief Executive Officer of Lancium.

~~20-19.~~ 20-19. On information and belief, Defendant Cline is the Chief Computing Officer of Lancium. On information and belief, as the Chief Computing Officer, Cline participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

~~21-20.~~ 21-20. Cline is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from Lancium's interest and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claim against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant Cline together. Plaintiffs' claims against Defendant Cline arise out of his exercise of his powers as Chief Computing Officer of Lancium.

~~22-21.~~ 22-21. The actions of Defendants McNamara and Cline establish sufficient minimum contacts with Delaware under Delaware law and the United States Constitution to give this Court personal jurisdiction over each of them.

~~23-22.~~ 23-22. As described below, each Defendant has committed acts giving rise to this action.

VENUE

24-23. Venue is proper in this District under 28 U.S.C. § 1391(b)(3) because there is no district in which an action may otherwise be brought as provided in § 1391(b) and Defendant Lancium is subject to the Court’s personal jurisdiction with respect to this action.

PLAINTIFFS’ PROPRIETARY CRYPTOCURRENCY MINING TECHNOLOGY

25-24. As of 2018, the amount of energy required to process computer algorithms to mine cryptocurrencies like Bitcoin was three times greater than the energy required to physically mine gold. Conventional mining of “copper, gold, platinum, and rare earth oxides are 4, 5, 7, and 9 megajoules to generate one U.S. dollars,” while “it costs an average of 17 megajoules to mine \$1 worth of bitcoin.”¹ The large amount of energy required to mine cryptocurrencies can make such mining financially prohibitive, and even when financially lucrative, the large energy requirements make cryptocurrency mining harmful to the global environment, with studies showing carbon dioxide emissions from cryptocurrency mining “single-handedly rais[ing] global temperatures by 2 degrees by 2023.” *Id.*

26-25. At the same time, some forms of electrical power generation are terribly inefficient. When producers of electrical power are unable to quickly adjust their operations in response to dynamically changing grid conditions, these producers frequently sell power at low or even negative prices until demand and market prices increase.

27-26. Because cryptocurrency mining is a computationally demanding process, it requires significant energy. As a result, industrial-scale cryptocurrency mining places a large

¹ <https://www.marketwatch.com/story/mining-bitcoin-is-3-times-more-expensive-than-mining-gold-research-paper-finds-2018-11-06>

energy burden on the power grid, driving demand and costs as well as increasing the likelihood of grid component failure.

28-27. In late 2018 and early 2019, Austin Storms sought to address these problems by developing energy-efficient cryptocurrency mining systems and methods that reduce the environmental impact of energy-intensive mining operations. Storms conceived of a system that better uses available energy resources to increase the stability of the energy grid, minimize a mining operation's impact on peak-demand, and alleviate ~~electricity undersupply and/or oversupply~~energy over supply conditions, all while decreasing the overall energy costs of the mining operation and increasing its profitability.

29-28. Austin Storms conceived of and developed the BearBox Technology. Storms is the president and founder of BearBox. The BearBox Technology includes hardware and software components. Structurally, the BearBox Technology includes a housing for a plurality of miners (such as ASICs, graphics cards, or the like) under the direction of a smart controller(s).

30-29. The smart controller monitors various external factors, such as current and expected energy demand and pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the system may determine whether conditions are appropriate to mine cryptocurrency and, if so, subsequently mines the cryptocurrency. Optionally, the system also includes other components for cooling, air-filtration, and related features.

31-30. In the BearBox Technology, a controller (such as a power distribution unit, network interface, or the like) monitors various external factors, such as current and expected energy demand/pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the controller(s) determines appropriate times to mine

cryptocurrency in accordance with a desired performance strategy (for example, profitability thresholds). At the appropriate times, the controller initiates mining, for example, by powering on the miners.

DEFENDANTS WRONGFULLY CLAIM THE BEARBOX TECHNOLOGY AS THEIR OWN

32.31. In May 2019, Storms attended the Fidelity FCAT Mining Summit in Boston, Massachusetts on behalf of BearBox to promote the BearBox Technology and seek potential customers for his revolutionary system.

33.32. While at the conference, Storms met Defendant McNamara. Defendant McNamara showed immediate interest in the BearBox Technology. Under the rouse of a potential business relationship, McNamara pumped Storms for details about the BearBox Technology over the course of several exchanges, which included conversations, emails, and text messages about the BearBox Technology. Storms took McNamara to dinner where McNamara continued to pump Storms for details about the BearBox Technology. At all times before and during Storms's disclosure of this information, Storms told McNamara that the BearBox Technology was confidential, and Storms relied on McNamara's good faith assurances that he would keep confidential the information he received from Storms about the BearBox Technology.

34.33. Following the conference, McNamara continued to press Storms for additional details about the BearBox Technology via text messaging and email. Again relying on Defendant McNamara's assurances of confidentiality, Storms provided annotated system diagrams, component specifications, and modeled data sets to mimic real-world Bitcoin and energy prices. Storms included express confidentiality notices in his communications with Defendant McNamara.

~~35.34.~~ After Storms disclosed the BearBox Technology to McNamara, McNamara abruptly ended all communications with Storms.

~~36.35.~~ Storms last communicated with McNamara on May 9, 2019 via e-mail, and after sending that message, Storms did not hear from McNamara again.

~~37.36.~~ At that time, Storms understood that McNamara was not interested in investing in the BearBox Technology. He had no reason to suspect that McNamara would steal the BearBox Technology and claim it as his own.

~~38.37.~~ On information and belief, Defendants filed U.S. provisional patent application No. 62/927,119 on October 28, 2019, naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

~~39.38.~~ In addition to falsely claiming to be the inventors of the inventions disclosed in the application, Defendants wrongfully disclosed, without authorization, the confidential BearBox Technology to the United States Patent and Trademark Office.

~~40.39.~~ Likewise, on December 4, 2019, Defendants filed U.S. Patent Application Serial No. 16/702,931, once again naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

~~41.40.~~ The '433 Patent issued on March 31, 2020 naming Defendants McNamara and Cline as the sole purported inventors on the face of the patent. A true and correct copy of the '433 Patent is attached hereto as Exhibit A.

~~42.41.~~ The inventions claimed in the '433 patent ~~encompass~~fall within the scope of the BearBox Technology, yet Defendants falsely identified themselves as the inventors of the claimed inventions, when, in fact, Storms is the sole inventor of the claimed inventions.

~~43.42.~~ On information and belief, McNamara and Cline assigned their purported rights in the '433 patent to Lancium. On information and belief, at all times, Lancium was aware that McNamara and Cline, both officers of Lancium, were not the rightful inventors of the BearBox Technology disclosed in the patent and the inventions claimed in the patent.

~~44.43.~~ Defendants McNamara and Cline each submitted signed declarations falsely swearing that they were “an original joint inventor” of the claimed subject matter. A true and correct copy of Defendant McNamara’s and Defendant Cline’s declarations are attached as Exhibit B.

~~45.44.~~ On August 14, 2020, Lancium filed a lawsuit in the U.S. District Court for the Western District of Texas against Layer1 Technologies, Inc. (“Layer1”) asserting that Layer1 infringes the '433 patent. That case is captioned *Lancium LLC v. Layer1 Technologies, Inc.*, Case No. 6:20-cv-739 (W.D. Texas) (the “Layer1 Lawsuit”).

~~46.45.~~ As part of the Layer1 Lawsuit, Defendants falsely asserted that McNamara and Cline are the sole inventors of the inventions claimed in the '433 patent.

~~47.46.~~ Plaintiffs became aware of Defendants’ wrongful use of the BearBox Technology on or about August 17, 2020, when they learned about the Layer1 Lawsuit through a press release dated August 14, 2020, posted by Lancium on PRNewswire. That press release is available at the following URL: <https://www.prnewswire.com/news-releases/controllable-load-resource-clr-market-leader-lancium-files-patent-infringement-lawsuit-against-layer1-301112687.html>.

~~48.47.~~ Before seeing the August 14, 2020 press release, Plaintiffs were unaware of Defendants’ wrongful use of the BearBox Technology and was unaware of the '433 patent.

~~49.48.~~ On March 5, 2021, Lancium and Layer1 entered a Stipulation to Dismiss with Prejudice in the Layer1 Lawsuit. According to the stipulation, the parties had entered a Settlement Agreement to resolve the Layer1 Lawsuit.

~~50.49.~~ According to a press release issued by Lancium on March 8, 2021, Lancium and Layer 1 “have entered into a mutually beneficial partnership. Layer1 has licensed Lancium’s intellectual property and Lancium will provide Smart Response™ software and services to Layer1.” The press release is available at the following URL: <https://www.prnewswire.com/news-releases/lancium-and-layer1-settle-patent-infringement-suit-301242602.html>

~~51.50.~~ On information and belief, as part of the Settlement Agreement between Lancium and Layer1 to settle the Layer1 Lawsuit, Lancium received and continues to receive valuable consideration from Layer1, all of which rightly belongs to Plaintiffs, the rightful owners of the inventions claimed in the ’433 Patent.

COUNT I
CORRECTION OF INVENTORSHIP FOR THE ’433 PATENT:
AUSTIN STORMS AS SOLE INVENTOR

~~52.51.~~ Plaintiffs incorporate the above paragraphs by reference.

~~53.52.~~ Storms is the sole inventor of the subject matter claimed in the ’433 Patent.

~~54.53.~~ Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the ’433 patent and the currently listed inventors on the ’433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms or BearBox.

~~55.54.~~ Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the ’433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT II
IN THE ALTERNATIVE, CORRECTION OF INVENTORSHIP FOR THE '433
PATENT: AUSTIN STORMS AS JOINT INVENTOR WITH THE CURRENTLY
NAMED INVENTORS

~~56.55.~~ Plaintiffs incorporates the above paragraphs by reference.

~~57.56.~~ In the alternative, Storms is a joint inventor of the subject matter claimed in the '433 Patent and should be added to the individuals currently named as inventors on the '433 Patent.

~~58.57.~~ Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the '433 patent and the currently listed inventors on the '433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms.

~~59.58.~~ Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the '433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT III
CONVERSION BY LANCIUM, MCNAMARA, AND CLINE

~~60.59.~~ Plaintiffs incorporates the above paragraphs by reference.

~~61.60.~~ Austin Storms, in his capacity as founder and President of BearBox, conceived, developed, and reduced to practice the BearBox Technology. Plaintiffs own the BearBox Technology, related know-how, and related intellectual property. Plaintiffs owned this property during all relevant time periods in this suit. Information on the BearBox Technology was provided to Defendants solely for the purposes of evaluation for a potential business relationship and under strict confidentiality obligations.

~~62-61.~~ Defendants assumed dominion and control over the BearBox Technology by claiming it as their own in the '433 patent. Through their wrongful conduct in obtaining the '433 Patent and claiming the BearBox Technology as their own, the Defendants have wrongfully obtained the purported ability to exclude Plaintiffs and others from using the BearBox Technology. This constitutes unauthorized and unlawful conversion by Defendants.

~~63-62.~~ As a result of Defendants' wrongful actions, Plaintiffs will suffer imminent and irreparable damages in an amount to be proven at trial. In particular, Plaintiffs have been damaged by losing valuable intellectual property from which Plaintiffs would have derived substantial revenue via licensing and/or selling patented products.

COUNT IV UNJUST ENRICHMENT BY LANCIUM, MCNAMARA, AND CLINE

~~64-63.~~ Plaintiffs incorporate the above paragraphs by reference.

~~65-64.~~ Plaintiffs conferred a benefit on Defendants by providing them valuable intellectual property about cryptocurrency mining systems and related confidential information and materials under the boundaries of a potential collaboration between BearBox and Lancium.

~~66-65.~~ Defendants accepted that cryptocurrency mining intellectual property and, indeed, continuously asked Storms to provide more information and materials, having recognized the benefit that Defendants received by having access to the BearBox Technology.

~~67-66.~~ Defendants accepted and retained the BearBox Technology, and used it to their own advantage, at Plaintiffs' expense.

~~68-67.~~ Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct. In particular, Defendants have unlawfully used Plaintiffs' property by asserting inventorship over the BearBox Technology, and deriving an unjust benefit from

exploiting Storms's cryptocurrency mining inventions. It would be inequitable for Defendants to retain these benefits under these circumstances.

~~69-68.~~ Plaintiffs have incurred, and continue to incur, detriment in the form of loss of money and property as a result of Defendants' wrongful use of Plaintiffs' intellectual property, including the right to any patent based on their own intellectual property ~~and any royalties derived from that intellectual property.~~ The intellectual property, including the right to any patents based on Plaintiffs' intellectual property and to any patent documents (including assignment documents), U.S. and foreign, are unique and there is no adequate remedy at law.

~~70-69.~~ The harm to Plaintiffs is continuous, substantial, and irreparable.

COUNT V
TRADE SECRET MISAPPROPRIATION UNDER
FEDERAL DEFEND TRADE SECRETS ACT, 18 U.S.C. § 1836

~~71. Plaintiffs incorporate the above paragraphs by reference.~~

~~72. At all relevant times, Plaintiffs and Defendants were engaged in interstate commerce. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods in interstate commerce throughout the United States.~~

~~73. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about the BearBox Technology relating to cryptocurrency mining systems and related methods from Plaintiff. The confidential information included detailed technical information about the BearBox Technology. Such information constitutes trade secrets of substantial economic value. This confidential technical information was not known to the public or to other persons who can obtain economic value from its~~

~~disclosure or use. This information constituted a trade secret under the Defend Trade Secrets Act, 18 U.S.C. §§ 1836.~~

~~74. Defendants knew the technical information they received from Austin Storms about the BearBox Technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, for internal use only, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship. Plaintiffs took reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its offices and computers.~~

~~75. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent constituted a misappropriation of Plaintiffs' trade secrets. Defendants acquired the confidential information under circumstances giving rise to a duty to maintain the secrecy of the trade secret or limit the use of the trade secret.~~

~~76. Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent was an unauthorized disclosure in breach of Defendants' agreement to maintain the secrecy of this information or otherwise limit its use. As recently as December 4, 2019, Defendants have wrongly sworn that Plaintiffs' confidential technical information is their own by declaring to the United States Patent and Trademark Office that they were the original and joint inventors of the inventions claimed in the '433 patent. Even ignoring that Defendants have improperly claimed this information as their own, through their public filings Defendants have disclosed the confidential technical information without the express or implied consent of~~

~~Plaintiffs. The acts of taking Plaintiffs' confidential information, claiming it as Defendants' own, and submitting it in public filings were improper means in breach of a confidential relationship. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.~~

~~77. — As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, damages for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability for a reasonable royalty for the Defendants' unauthorized disclosure or use of the trade secrets.~~

~~78. — The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D).~~

~~COUNT VI~~

~~TRADE SECRET MISAPPROPRIATION UNDER~~

~~TEXAS CIVIL PRACTICE AND REMEDIES CODE § 134A~~

~~79. — Plaintiffs incorporate the above paragraphs by reference.~~

~~80. — At all relevant times, Plaintiffs and Defendants were engaged in trade or commerce within the meaning of Texas Civil Practice and Remedies c. § 134A, the Texas Uniform Trade Secrets Act. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods.~~

81. — In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about the BearBox Technology relating to cryptocurrency mining systems and related methods from Plaintiff. The confidential information included detailed technical information about the BearBox Technology. Such information constitutes trade secrets of substantial economic value. This confidential technical information was not known to the public or to other persons who can obtain economic value from its disclosure or use. This information constituted a trade secret.

82. — Defendants knew the technical information they received from Austin Storms about the BearBox Technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, for internal use only, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship. Plaintiffs took reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its offices and computers.

83. — Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent constituted a misappropriation of Plaintiffs' trade secrets. Defendants acquired the confidential information under circumstances giving rise to a duty to maintain the secrecy of the trade secret or limit the use of the trade secret.

84. — Defendants' inclusion of Plaintiffs' confidential technical information in the '433 patent was an unauthorized disclosure in breach of Defendants' agreement to maintain the secrecy of this information or otherwise limit its use. As recently as December 4, 2019,

~~Defendants have wrongly sworn that Plaintiffs' confidential technical information is their own by declaring to the United States Patent and Trademark Office that they were the original and joint inventors of the inventions claimed in the '433 patent. Even ignoring that Defendants have improperly claimed this information as their own, through their public filings Defendants have disclosed the confidential technical information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, claiming it as Defendants' own, and submitting it in public filings were improper means in breach of a confidential relationship. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.~~

~~85. — As a result of the misappropriation of Plaintiffs' trade secrets and intended immediate use by the Defendants, Plaintiffs' have suffered and will suffer damages in an amount to be proven at trial.~~

~~COUNT VI~~
COUNT V
NEGLIGENT MISREPRESENTATION BY LANCIMUM AND MCNAMARA

~~86.70.~~ Plaintiffs incorporate the above paragraphs by reference.

~~87.71.~~ In connection with the potential work involving cryptocurrency mining systems and related methods, Storms told Defendant McNamara that the cryptocurrency mining systems and related methods were proprietary to Plaintiffs and not to be used or shared outside of Lancium. Defendant McNamara gave his word that he would abide by this confidentiality. On information and belief, Defendant McNamara agreed to keep the BearBox Technology confidential despite later recklessly incorporating the BearBox Technology into his own patent applications and swearing, as recently as December 4, 2019, that he is an inventor of the

BearBox Technology. Storms relied on Defendant McNamara's assurances of confidentiality and continued to share details about the BearBox Technology with Defendants.

~~88.72.~~ If Plaintiffs had known that Defendants would secretly incorporate the BearBox Technology into Defendants' own patent applications to claim them as Defendants' intellectual property, Plaintiffs would not have continued working with and sharing intellectual property with Defendants.

~~89.73.~~ Plaintiffs suffered a pecuniary loss based on this reliance including the loss of potential patent rights, and the costs of Plaintiffs' know-how converted under the guise of a potential business relationship.

JURY DEMAND

~~90.74.~~ Under Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiffs respectfully demand a trial by jury on all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, BearBox respectfully requests the following relief:

A. An order that the Director of the United States Patent and Trademark Office correct the inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

B. Alternatively, an order that Defendants sign the requisite documents to correct inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

C. A declaration that Austin Storms is the sole inventor, or, in the alternative, is a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

D. A preliminary and a permanent injunction enjoining Defendants Lancium, McNamara, and Cline from asserting that McNamara or Cline are inventors of the '433 Patent in violation of the United States federal patent laws;

E. An order that Defendants immediately transfer to Plaintiffs all right, title, and interest in all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

F. An order for a constructive trust over all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of Plaintiffs' intellectual property;

G. Financial relief including damages, consequential damages, disgorgement of Defendants' ill-gotten profits, Defendants' unjust enrichment, ~~restitution~~, reasonable royalty damages, lost profits damages, reliance damages, and/or all other appropriate financial relief, all in an amount to be determined at trial, with interest;

H. An award of the amount by which Defendants have been unjustly enriched by their actions set forth in this Complaint and their purported ownership of patents covering Plaintiffs' intellectual property;

~~I. An award of exemplary, enhanced, and/or punitive damages as permitted by law, including under 18 U.S.C. § 1836(b)(3)(C) and Texas Business and Commerce Code § 17E;~~

I. A finding that this is an exceptional case warranting imposition of attorney fees against Defendants and an award to Plaintiffs of its reasonable costs and attorney fees incurred in bringing this action pursuant to 35 U.S.C. § 285; and

~~K. An award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(D); and~~

L.J. An award of such further relief at law or in equity, such as preliminary and/or permanent injunctive relief, as the Court deems just and proper.

ASHBY & GEDDES

/s/ Andrew C. Mayo

Andrew C. Mayo (#5207)
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19899
(302) 654-1888
amayo@ashbygeddes.com

Of Counsel:

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati, III
Chelsea M. Murray
MARSHALL, GERSTEIN & BORUN LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
(312) 474-6300

*Attorneys for Plaintiffs
BearBox LLC and Austin Storms*

Dated: May 24, 2021

LANCIUM IS THE TRUE INNOVATOR IN THIS FIELD

20. Lancium is a technology company creating software and intellectual property solutions that enable more renewable energy on the nation's power grid. Lancium's products include Lancium Smart Response™ for rapid server power management, and Lancium Compute™, a platform for high throughput computing applications. Lancium's solutions help ensure that renewable energy can power our future.

21. Lancium is the leader in data center power ramping software.

22. As data center power demand continues to escalate, Lancium Smart Response™ software can unlock huge power costs savings for data center owners. The software also provides critical services to the power grid ensuring reliability and resiliency. As every grid takes on more renewable energy, the power market will need much greater quantities of flexible load, and in particular, Controllable Loads enabled by Lancium Smart Response™. Lancium Smart Response™ functionality enables data centers to provide this crucial service.

23. Lancium has numerous issued and pending patents.

24. Lancium's issued patents include the following:

- U.S. Patent No. 10,873,211 (filed Sept. 13, 2018) "Systems and Methods for Dynamic Power Routing with Behind-the-Meter Energy Storage"
- U.S. Patent No. 10,444,818 (filed Oct. 30, 2018) "Methods and Systems for Distributed Power Control of Flexible Datacenters"
- U.S. Patent No. 10,367,353 (filed Oct. 30, 2018) "Managing Queue Distribution between Critical Datacenter and Flexible Datacenter"
- U.S. Patent No. 10,452,127 (filed Jan. 11, 2019) "Redundant Flexible Datacenter Workload Scheduling"
- U.S. Patent No. 10,618,427 (filed Oct. 8, 2019) "Behind-the-Meter Branch Loads for Electrical Vehicle Charging"
- U.S. Patent No. 10,608,433 (filed Dec. 4, 2019) "Methods and Systems for Adjusting Power Consumption Based on a Fixed-Duration Power Option"

Agreement”

- U.S. Patent No. 10,857,899 (filed Mar. 4, 2020) “Behind-the-Meter Branch Loads for Electrical Vehicle Charging”

25. One of Lancium’s patent applications, International Application Number PCT/US2018/0 17950, was filed on February 13, 2018 and published as the ‘632 application on July 18, 2019.

26. The ‘632 application describes a method and system for dynamic power delivery to a flexible datacenter using unutilized energy sources. In one embodiment of the invention:

a method and system for dynamic power delivery' to a flexible datacenter uses unutilized behind-the-meter power sources without transmission and distribution costs. The flexible datacenter may be configured to modulate power delivery to one or more computing systems based on the availability of unutilized behind-the-meter power or an operational directive. For example, the flexible datacenter may ramp-up to a fully online status, ramp-down to a fully offline status, or dynamically reduce power consumption, act a load balancer, or adjust the power factor. Advantageously, the flexible datacenter may perform computational operations, such as blockchain hashing operations, with little to no energy costs, using clean and renewable energy that would otherwise be wasted.

See ‘632 application ¶ [0022].

27. Put simply, and as discussed more thoroughly below, the ‘632 application describes the so-called “BearBox technology.”

28. The ‘632 application claims priority to U.S. Provisional Application No. 62/616,348 (the “’348 Provisional”), filed on Jan. 11, 2018.

29. The ‘632 application lists the following Lancium current or former personnel as inventors: David Henson, Michael McNamara, and Raymond Cline.

30. The ‘632 application lists the “Applicant” as Lancium LLC.

31. On information and belief, BearBox LLC knew of the ‘632 application before filing the Complaint in this action.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,

Plaintiffs,

v.

LANCIUM LLC, MICHAEL T.
MCNAMARA, and RAYMOND E.
CLINE, JR.,

Defendants.

C.A. No. 21-534-MN

SCHEDULING ORDER

This 6th day of July 2021, ~~the Court having conducted an initial Rule 16(b) scheduling conference pursuant to Local Rule 16.1(b), and~~ the parties having determined after discussion that the matter cannot be resolved at this juncture by settlement, voluntary mediation, or binding arbitration;

IT IS ORDERED that:

1. Rule 26(a)(1) Initial Disclosures and E-Discovery Default Standard.

Unless otherwise agreed to by the parties, the parties shall make their initial disclosures pursuant to Federal Rule of Civil Procedure 26(a)(1) within five (5) days of the date the Court entered this Order. If they have not already done so, the parties are to review the Court's Default Standard for Discovery, Including Discovery of Electronically Stored Information ("ESI"), which is posted at <http://www.ded.uscourts.gov> (see Other Resources, Default Standard for Discovery) and is incorporated herein by reference.

2. Joinder of Other Parties and Amendment of Pleadings. All motions to join other parties, and to amend or supplement the pleadings, shall be filed on or before **November 1, 2021**.

Unless otherwise ordered by the Court, any motion to join a party or motion to amend the pleadings shall be made pursuant to the procedures set forth in Paragraphs 7(g) and 8.

3. Application to Court for Protective Order. Should counsel find it will be necessary to apply to the Court for a protective order specifying terms and conditions for the disclosure of confidential information, counsel should confer and attempt to reach an agreement on a proposed form of order and submit it to the Court within ten (10) days from the date the Court enters this Order. Should counsel be unable to reach an agreement on a proposed form of order, counsel must follow the provisions of Paragraph 7(g) below.

Any proposed protective order must include the following paragraph:

Other Proceedings. By entering this order and limiting the disclosure of information in this case, the Court does not intend to preclude another court from finding that information may be relevant and subject to disclosure in another case. Any person or party subject to this order who becomes subject to a motion to disclose another party's information designated "confidential" [the parties should list any other level of designation, such as "highly confidential," which may be provided for in the protective order] pursuant to this order shall promptly notify that party of the motion so that the party may have an opportunity to appear and be heard on whether that information should be disclosed.

4. Papers Filed Under Seal. In accordance with section G of the Revised Administrative Procedures Governing Filing and Service by Electronic Means, a redacted version of any sealed document shall be filed electronically within seven (7) days of the filing of the sealed document.

5. Courtesy Copies. The parties shall provide to the Court two (2) courtesy copies of all briefs and any other document filed in support of any briefs (*i.e.*, appendices, exhibits, declarations, affidavits etc.). This provision also applies to papers filed under seal. All courtesy copies shall be double-sided.

6. ADR Process. This matter is referred to a magistrate judge to explore the possibility of alternative dispute resolution.

7. Discovery. Unless otherwise ordered by the Court or agreed to by parties, the limitations on discovery set forth in the Federal Rules shall be strictly observed.

(a) Fact Discovery Cut Off. All fact discovery in this case shall be initiated so that it will be completed on or before December 14, 2021.

(b) Document Production. Document production shall be substantially complete by October 8, 2021.

(c) Requests for Admission. A maximum of fifty (50) requests for admission are permitted for each side. For clarity, this limit does not apply to requests for admission related solely to authenticity.

(d) Interrogatories.

i. A maximum of twenty-five (25) interrogatories, including contention interrogatories, are permitted for each side.

ii. The Court encourages the parties to serve and respond to contention interrogatories early in the case. In the absence of agreement among the parties, contention interrogatories, if filed, shall first be addressed by the party with the burden of proof. The adequacy of all interrogatory answers shall be judged by the level of detail each party provides (*i.e.*, the more detail a party provides, the more detail a party shall receive).

(e) Depositions.

i. Limitation on Hours for Deposition Discovery. Each side is limited to a total of thirty (30) hours of taking testimony by deposition upon oral examination.

ii. Location of Depositions. Any party or representative (officer, director, or managing agent) of a party filing a civil action in this district court must ordinarily be required, upon request, to submit to a deposition at a place designated within this district. Exceptions to this general rule may be made by order of the Court. A defendant who becomes a counterclaimant, cross-claimant, or third-party plaintiff shall be considered as having filed an action in this Court for the purpose of this provision.

(f) Disclosure of Expert Testimony.

i. Expert Reports. For the party who has the initial burden of proof on the subject matter, the initial Federal Rule of Civil Procedure 26(a)(2) disclosure of expert testimony is due on or before January 18, 2022. The supplemental disclosure to contradict or rebut evidence on the same matter identified by another party is due on or before February 18, 2022. Reply expert reports from the party with the initial burden of proof are due on or before March 8, 2022. No other expert reports will be permitted without either the consent of all parties or leave of the Court. Along with the submissions of the expert reports, the parties shall advise of the dates and times of their experts' availability for deposition.

ii. Objections to Expert Testimony. To the extent any objection to expert testimony is made pursuant to the principles announced in *Daubert v. Merrell Dow Pharmaceuticals, Inc.*, 509 U.S. 579 (1993), as incorporated in Federal Rule of Evidence 702, it shall be made by motion no later than the deadline for dispositive motions set forth herein, unless otherwise ordered by the Court.

iii. Expert Discovery Cut Off. All expert discovery in this case shall be initiated so that it will be completed on or before April 13, 2022.

(g) Discovery Matters and Disputes Relating to Protective Orders.

i. Any discovery motion filed without first complying with the following procedures will be denied without prejudice to renew pursuant to these procedures.

ii. Should counsel find, after a reasonable effort pursuant to Local Rule 7.1.1 that they are unable to resolve a discovery matter or a dispute relating to a protective order, the parties involved in the discovery matter or protective order dispute shall contact the Court's Judicial Administrator to schedule an argument.

iii. On a date to be set by separate order, generally not less than four (4) days prior to the conference, the party seeking relief shall file with the Court a letter, not to exceed three (3) pages, outlining the issues in dispute and its position on those issues. On a date to be set by separate order, but generally not less than three (3) days prior to the conference, any party opposing the application for relief may file a letter, not to exceed three (3) pages, outlining that party's reasons for its opposition.

iv. The parties shall provide to the Court two (2) courtesy copies of its discovery letter and any other document filed in support of any letter (*i.e.*, appendices, exhibits, declarations, affidavits etc.). This provision also applies to papers filed under seal. All courtesy copies shall be double-sided.

v. Should the Court find further briefing necessary upon conclusion of the conference, the Court will order it. Alternatively, the Court may choose to resolve the dispute prior to the conference and will, in that event, cancel the conference.

8. Claim Construction. The parties have determined that claim construction may be necessary in this case. By August 27, 2021, the parties shall determine whether or not claim construction is necessary and further notify the Court of their decision, and if claim construction is necessary, the parties will propose a claim construction schedule for the Court's consideration.

9. Motions to Amend / Motions to Strike.

(a) Any motion to amend (including a motion for leave to amend) a pleading or any motion to strike any pleading or other document shall be made pursuant to the discovery dispute procedure set forth in Paragraph 7(g) above.

(b) Any such motion shall attach the proposed amended pleading as well as a “redline” comparison to the prior pleading or attach the document to be stricken.

10. Case Dispositive Motions.

(a) All case dispositive motions, an opening brief, and affidavits, if any, in support of the motion shall be served and filed on or before June 13, 2022. Briefing will be presented pursuant to the Court’s Local Rules. No case dispositive motion under Rule 56 may be filed more than ten (10) days before the above date without leave of the Court.

(b) Concise Statement of Facts Requirement. Any motion for summary judgment shall be accompanied by a separate concise statement, not to exceed six (6) pages, which details each material fact which the moving party contends is essential for the Court’s resolution of the summary judgment motion (not the entire case) and as to which the moving party contends there is no genuine issue to be tried. Each fact shall be set forth in a separate numbered paragraph and shall be supported by specific citation(s) to the record.

Any party opposing the motion shall include with its opposing papers a response to the moving party’s concise statement, not to exceed six (6) pages, which admits or disputes the facts set forth in the moving party’s concise statement on a paragraph-by-paragraph basis. To the extent a fact is disputed, the basis of the dispute shall be supported by specific citation(s) to the record. Failure to respond to a fact presented in the moving party’s concise statement of facts shall indicate that fact is not in dispute for purposes of summary judgment. The party opposing the

motion may also include with its opposing papers a separate concise statement, not to exceed four (4) pages, which sets forth material facts as to which the opposing party contends there is a genuine issue to be tried. Each fact asserted by the opposing party shall also be set forth in a separate numbered paragraph and shall be supported by specific citation(s) to the record.

The moving party shall include with its reply papers a response to the opposing party's concise statement of facts, not to exceed four (4) pages, on a paragraph-by-paragraph basis. Failure to respond to a fact presented in the opposing party's concise statement of facts shall indicate that fact remains in dispute for purposes of summary judgment.

11. Applications by Motion. Except as otherwise specified herein, any application to the Court shall be by written motion. Any non-dispositive motion should contain the statement required by Local Rule 7.1.1.

12. Motions in Limine. Motions *in limine* shall not be separately filed. All *in limine* requests and responses thereto shall be set forth in the proposed pretrial order. Each party shall be limited to three (3) *in limine* requests, unless otherwise permitted by the Court. The *in limine* request and any response shall contain the authorities relied upon; each *in limine* request may be supported by a maximum of three (3) pages of argument, may be opposed by a maximum of three (3) pages of argument, and the party making the *in limine* request may add a maximum of one (1) additional page in reply in support of its request. If more than one party is supporting or opposing an *in limine* request, such support or opposition shall be combined in a single three (3) page submission (and, if the moving party, a single one (1) page reply), unless otherwise ordered by the Court. No separate briefing shall be submitted on *in limine* requests, unless otherwise permitted by the Court.

13. Pretrial Conference. On **November 22, 2022**, the Court will hold a pretrial conference in Court with counsel beginning at **4:30 p.m.** The parties shall file with the Court the joint proposed final pretrial order in compliance with Local Rule 16.3(c) and the Court's Preferences and Procedures for Civil Cases not later than seven (7) days before the pretrial conference. Unless otherwise ordered by the Court, the parties shall comply with the timeframes set forth in Local Rule 16.3(d)(1)-(3) for the preparation of the joint proposed final pretrial order. The Court will advise the parties at or before the above-scheduled pretrial conference whether an additional pretrial conference will be necessary.

The parties shall provide the Court two (2) double-sided courtesy copies of the joint proposed final pretrial order and all attachments. The proposed final pretrial order shall contain a table of contents and the paragraphs shall be numbered.

14. Jury Instructions, Voir Dire, and Special Verdict Forms. Where a case is to be tried to a jury, pursuant to Local Rules 47.1(a)(2) and 51.1 the parties should file (i) proposed voir dire, (ii) preliminary jury instructions, (iii) final jury instructions, and (iv) special verdict forms seven (7) full business days before the final pretrial conference. This submission shall be accompanied by a courtesy copy containing electronic files of these documents, in Microsoft Word format, which may be submitted by e-mail to mn_civil@ded.uscourts.gov.

15. Trial. This matter is scheduled for a four (4) day jury trial beginning at 9:30 a.m. on **December 5, 2022**, with the subsequent trial days beginning at 9:00 a.m. Until the case is submitted to the jury for deliberations, the jury will be excused each day at 4:30 p.m. The trial will be timed, as counsel will be allocated a total number of hours in which to present their respective cases.



The Honorable Maryellen Noreika
United State District Judge

EXHIBIT A

EVENT	DEADLINE
Submission of Joint Scheduling Order	June 24, 2021
Rule 26(a)(1) Initial Disclosures	Within 5 days of the date of the Scheduling Order
Application to Court for Protective Order	Within 10 days of the date of the Scheduling Order
Determination of Whether Claim Construction is Necessary and Advise Court	August 27, 2021
Substantial Completion of Document Production	October 8, 2021
Fact Discovery Cut-Off	December 14, 2021
Initial Expert Reports	January 18, 2022
Rebuttal Expert Reports	February 22, 2022
Reply Expert Reports	March 8, 2022
Expert Discovery Cut-Off	April 13, 2022
Case Dispositive Motions	June 13, 2022
Pretrial Conference	November 22, 2022 at 4:30 p.m.
Trial (4-day jury)	December 5, 2022

favorable to Plaintiffs, Plaintiffs have pled valid facts to support a finding that Defendants owed Plaintiffs a duty and Defendants' motion must be denied.

C. Material factual disputes preclude judgment on the pleadings

Defendants' Motion for Judgment on the Pleadings should be denied because several material issues of fact exist that, at a minimum, prevent resolution of necessary issues at this time. Judgment on the pleadings is proper only if the "movant clearly establishes that no material issue of fact remains to be resolved." *Mid-Am. Salt, LLC v. Morris Cty. Coop. Pricing Council*, 964 F.3d 218, 226 (3d Cir. 2020). Yet Defendants do not even attempt to argue that there are no issues of material fact.

A cursory review of the pleadings show that this case is replete with material issues of fact that preclude judgment on the pleadings. "[D]efendants filed an answer, denied liability, and have asserted affirmative defenses. The defendants' denials are sufficient to place *all* material facts in this case in dispute and, therefore, serve as a bar to [defendants'] motion for judgment on the pleadings." *Rainier v. Rispoli*, 2012 WL 752371, at *2 (D. Del. 2012) (emphasis added). Defendants' answer here, particularly when viewed in favor of Plaintiffs, contain denials of liability and asserted affirmative defenses that place these material facts at issue in this case and thereby bar the instant motion. *See, e.g.*, Amended Answer to Amended Complaint, D.I. 30, ¶¶ 1, 3-4, 27-50; *Rainier*, 2012 WL 752371, at *2 (denying motion for judgment on the pleadings).

As an example, there are several material facts in dispute that are relevant to Plaintiffs' conversion claim in light of Defendants' Amended Answer. Under Louisiana law, "[t]o prevail on a claim of conversion, a plaintiff must prove that (1) he owned or had the right to possess; (2) the defendant's use was inconsistent with the plaintiff's right of ownership; and (3) the defendant's use constituted a wrongful taking." *Mabile v. BP, p.l.c.*, No. CV 11-1783, 2016 WL

ASHBY & GEDDES

ATTORNEYS AND COUNSELLORS AT LAW

500 DELAWARE AVENUE

P. O. BOX 1150

WILMINGTON, DELAWARE 19899

TELEPHONE
302-654-1888

FACSIMILE
302-654-2067

August 31, 2021

The Honorable Maryellen Noreika
United States District Court
844 N. King Street
Wilmington, DE 19801

VIA ELECTRONIC FILING

Re: *Bear Box LLC, et al. v. Lancium LLC, et al.*,
C.A. No. 21-534-MN

Dear Judge Noreika:

I write on behalf of the parties in connection with Paragraph 8 the Scheduling Order (D.I. 35) regarding the need for a claim construction hearing in this matter. Plaintiffs have not yet made a determination regarding a need for a claim construction hearing, but anticipate any request for a claim construction hearing will be made prior to October 15, 2021 in view of the current progress of fact discovery (including contention discovery).¹ Defendants do not believe claim construction will be needed, and Defendants oppose Plaintiffs' attempt to extend the deadline in Paragraph 8 of the Scheduling Order until October 15, 2021.

Should the Court have any questions, counsel are available at the Court's convenience.

Respectfully,

Andrew C. Mayo

Andrew C. Mayo (#5702)

ACM/nlm

cc: All Counsel of Record (via electronic mail)

¹ The parties are continuing to exchange contention discovery that Plaintiffs believe may bear on whether claim construction is needed. The Plaintiffs anticipate completing this contention discovery without the need for motion practice before Magistrate Judge Burke. Although Defendants declined to stipulate to extend this deadline, Plaintiffs will not burden the Court with a formal motion at this time. Instead, Plaintiffs will seek a claim construction hearing only if necessary after completing additional contention discovery.

BARNES & THORNBURG LLP

1000 N. West Street
Suite 1500
Wilmington, Delaware 19801
302-300-3434
302-300-3456 (Fax)
www.btlaw.com

Chad S.C. Stover
302-300-3474
Chad.Stover@btlaw.com

October 15, 2021

BY ECF

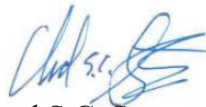
The Honorable U.S. Magistrate Judge Christopher J. Burke
U.S. District Court for the District of Delaware
J. Caleb Boggs Federal Building
844 N. King Street, Room 2325 – Unit 28
Wilmington, Delaware 19801

**Re: *BearBox LLC, et al. v. Lancium LLC, et al.*,
C.A. No. 21-534-MN-CJB**

Dear Judge Burke:

Pursuant to Court's September 13, 2021 Oral Order (D.I. 59), I write to state that Defendants do not request claim construction at this time but reserve their right to later seek claim construction if Plaintiffs later seek claim construction, or based on information obtained during discovery.

Respectfully submitted,



Chad S.C. Stover (DE No. 4919)

cc: Andrew Colin Mayo (via ECF)
Benjamin T. Horton (via ECF)
Chelsea M. Murray (via ECF)
John R. Labbe (via ECF)
Raymond R. Ricordati, III (via ECF)

Courts have found that claims for conversion⁴ that are dependent on a determination of patent inventorship are generally preempted by federal patent law. *See Gerawan Farming, Inc. v. Rehrig Pac. Co.*, No. 1:11-cv-01273-LJO BAM, 2012 WL 691758, at *7 (E.D. Cal. Mar. 2, 2012) (citing cases).⁵ Defendants assert that Plaintiffs' conversion claim requires a determination of inventorship and is therefore preempted. (D.I. 33 at 7-9; D.I. 46 at 2) The Court agrees.

Plaintiffs' Amended Complaint begins by alleging that the case "is about the Defendants' *theft of inventions that rightfully belong to Plaintiffs[.]*" and Plaintiffs' conversion claim begins by incorporating by reference the allegations from Plaintiffs' correction of inventorship claims. (D.I. 19 at ¶¶ 1, 59 (emphasis added)) From there, the conversion claim alleges that "Defendants assumed dominion and control over the BearBox Technology *by claiming it as their own in the '433 patent*" and that "[t]hrough their wrongful conduct *in obtaining the '433 [p]atent and claiming the BearBox Technology as their own*, [] Defendants have wrongfully obtained the purported ability to exclude Plaintiffs and others from using the BearBox Technology." (*Id.* at ¶ 61 (emphasis added)) The harm that Plaintiffs allege as a result of the conversion is the loss of "valuable intellectual property[.]" (*Id.* at ¶ 62) This is clearly a claim that, by its wording, is

⁴ At least for purposes of this Motion, both parties agree that Louisiana law applies to the state law claims at issue here. (D.I. 33 at 13; D.I. 42 at 10) Under Louisiana law, a claim for conversion requires proof of: (1) the plaintiff's right to possess; (2) that the defendant's use was inconsistent with the plaintiff's right of ownership; and (3) that the defendant's use constituted a wrongful taking. *Mabile v. BP, P.L.C.*, CIVIL ACTION NO. 11-1783, 2016 WL 5231839, at *21 (E.D. La. Sept. 22, 2016).

⁵ On the other hand, conversion claims that are based on a non-patent-ownership theory of conversion are generally not preempted by federal patent law. *Gerawan Farming, Inc.*, 2012 WL 691758, at *7.

dependent on a determination of patent inventorship. It is thus preempted by federal patent law.⁶ See, e.g., *Speedfit LLC v. Woodway USA, Inc.*, 226 F. Supp. 3d 149, 160 (E.D.N.Y. 2016) (concluding that the plaintiffs' conversion claim was preempted by federal patent law, where plaintiffs' "assertion of ownership over the subject matter of the Woodway Patents, which underlies their claim of conversion, clearly turns on a determination of inventorship"); *Gerawan Farming, Inc.*, 2012 WL 691758, at *7 (concluding that the plaintiff's conversion claim was preempted by federal patent law, where the claim alleged that the defendant "substantially interfered with [the plaintiff's] rightful property by omitting [the plaintiff] as an inventor" and accordingly "wrongfully exercised control over [the plaintiff's] rights in the . . . '293 Patent") (internal quotation marks and citation omitted); *Gen. Elec. Co. v. Wilkins*, No. 1:10-cv-00674-OWW-JLT, 2011 WL 3163348, at *9 (E.D. Cal. July 26, 2011) (finding that the defendant's conversion claim was preempted, where it sought "compensation for [p]laintiff's purported interference with [d]efendant's inventorship interests in the intellectual property embodied in the '565 and '985 patents; in other words, [d]efendant seeks patent-like protection under the guise of the tort of conversion").⁷

⁶ In their briefing, Plaintiffs try to recast the conversion claim as one "based on acts of misappropriating documents and information[.]" (D.I. 42 at 7) But, as shown above, that is not accurate. The actual allegations in Count III make clear that Plaintiffs' conversion claim is "based on the alleged right to ownership of the '433 patent due to Storms' alleged inventorship." (D.I. 46 at 4)

⁷ Moreover, the fact that Plaintiffs' conversion claim seeks damages amounting to the loss of "valuable intellectual property from which Plaintiffs would have derived substantial revenue via licensing and/or selling patented products[.]" (D.I. 19 at ¶ 62), demonstrates that it is patent-like in nature, see *Speedfit LLC*, 226 F. Supp. 3d at 160. That is another reason why the claim is preempted. (D.I. 46 at 4)

Because the Court finds that Plaintiffs' conversion claim is preempted, it need not now further evaluate Defendants' additional arguments for dismissal of the claim.

B. Unjust Enrichment (Count IV)

The next claim, in Count IV, is a claim for unjust enrichment.⁸ As with conversion claims, unjust enrichment claims that are based on a determination of patent inventorship are generally preempted by federal patent law. *See, e.g., Tavory v. NTP, Inc.*, 297 F. App'x 976, 983-84 (Fed. Cir. 2008) (finding that an unjust enrichment claim was preempted by federal patent law where the "dispositive issue [was plaintiff's] alleged co-inventorship"); *Heat Techs., Inc. v. Papierfabrik Aug. Koehler SE*, Civil Action No. 1:18-cv-01229-SDG, 2020 WL 12309512, at *4 (N.D. Ga. June 5, 2020) (citing cases); *James v. J2 Cloud Servs. Inc.*, Case No. 2:16-cv-05769-CAS(PJWx), 2018 WL 6092461, at *5 (C.D. Cal. Nov. 19, 2018).⁹ Defendants

⁸ The elements of an unjust enrichment claim in Louisiana are: (1) an enrichment of the defendant; (2) an impoverishment of the plaintiff; (3) a connection between the enrichment and the resulting impoverishment; (4) an absence of justification or cause for the enrichment and impoverishment; and (5) there must be no other remedy at law available to plaintiff. *Garber v. Badon & Ranier*, 981 So. 2d 92, 100 (La. Ct. App. 2008).

⁹ On the other hand, there are some circumstances where federal courts have concluded that an unjust enrichment claim having some connection to patent rights was not preempted, such as where it was alleged that defendants wrongfully obtained a patent based on confidential information received from a plaintiff, and claimed ownership of it in violation of a contract. *See, e.g., Spears v. SHK Consulting & Dev., Inc.*, 338 F. Supp. 3d 1272, 1277 (M.D. Fla. 2018); *see also Univ. of Colo. Found., Inc. v. Am. Cyanamid Co.*, 342 F.3d 1298, 1306-07 (Fed. Cir. 2003) (finding that an unjust enrichment claim was not preempted, where it amounted to "a legal claim to remedy the breach of a contract implied in law for disclosure of [plaintiffs'] confidential manuscript in exchange for a promise not to disseminate the idea without the [plaintiffs'] consent"); *cf. Ki Beom Kim v. Dyna Flex, Ltd.*, 525 F. Supp. 3d 999, 1007-08 (E.D. Mo. 2021) (denying a motion to remand, where plaintiff's unjust enrichment claim raises a substantial federal question regarding inventorship, as it was not "a dispute strictly concerning ownership, typically involving an underlying contract" but instead one where "Plaintiff is unquestionably arguing that he is the true inventor of the Invention and should have been included on the patents. If [Defendant] is the true inventor and owner of the patents, there is no reasonable claim that it has been unjustly enriched by continuing to sell products based on the Patents") (emphasis added).

assert that Plaintiffs’ unjust enrichment claim, like its conversion claim, is focused on the question of inventorship (and necessarily requires a determination of inventorship in order to know whether a legal violation can be found). (D.I. 33 at 9-11; D.I. 46 at 2-3) The Court again agrees with Defendants.

Plaintiffs’ unjust enrichment claim begins by incorporating the previous allegations of Plaintiffs’ inventorship and conversion claims. (D.I. 19 at ¶ 63) It then alleges that “Plaintiffs conferred a benefit on Defendants by providing them *valuable intellectual property* about cryptocurrency mining systems and related confidential information and materials” and that “Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct[.] *[i]n particular, Defendants have unlawfully used Plaintiffs’ property by asserting inventorship over the BearBox Technology, and deriving an unjust benefit from exploiting Storms’ cryptocurrency mining inventions.*” (*Id.* at ¶¶ 64, 67 (emphasis added)) As for damages, Plaintiffs allege a “loss of money and property as a result of Defendants’ *wrongful use of Plaintiffs’ intellectual property, including the right to any patent* based on their own intellectual property.” (*Id.* at ¶ 68 (emphasis added)) Thus, in light of the way Plaintiffs have pleaded the claim, the claim’s gravamen is that Defendants have been enriched because they have declared that they are the inventors of the '433 patent—and that this enrichment is unjust because Mr. Storms is actually the true inventor. This claim, as pleaded, is also preempted by federal patent law. *See James*, 2018 WL 6092461, at *4-5 (finding that the plaintiff’s unjust enrichment claim was preempted because it was “fundamentally based on plaintiff’s assertion that he should have been named on the '638 patent”); *OptoLum, Inc. v. Cree, Inc.*, 244 F. Supp. 3d 1005, 1014 (D. Ariz. 2017) (“Because OptoLum’s unjust enrichment claim depends on the

determination that OptoLum, not Cree, invented the LED technology at issue here, the claim is preempted by federal patent law.”).

Here again, because the Court finds that this claim is preempted, it declines to evaluate Defendants’ remaining arguments for dismissal.

C. Negligent Misrepresentation (Count V)

The Court turns lastly to Plaintiffs’ negligent misrepresentation claim in Count V. Under Louisiana law, the tort of negligent misrepresentation has three elements: (1) there must be a legal duty on the part of the defendant to supply correct information to the plaintiff; (2) there must be a breach of that duty; and (3) the breach must have caused damages to the plaintiff. *Schaumburg v. State Farm Mut. Auto. Ins. Co.*, 421 F. App’x 434, 439-40 (5th Cir. 2011); *Cook v. Am. Gateway Bank*, 49 So. 3d 23, 32 (La. Ct. App. 2010).

Defendants first assert that Count V is preempted because it is “dependent upon [Plaintiffs’] assertion that [Mr.] Storms is an inventor of the '433 patent[.]” (D.I. 33 at 11-12; *see also* D.I. 46 at 2-3) In the Court’s view, however, the argument for preemption here is not as strong as it was regarding Counts III and IV.

While Plaintiffs’ conversion and unjust enrichment claims clearly turned on whether Mr. Storms is a proper inventor of the '433 patent, the premise of Plaintiffs’ negligent misrepresentation claim is a bit different. In Count V, Plaintiffs first allege that Mr. Storms “told [Mr.] McNamara that the cryptocurrency mining systems and related methods were proprietary to Plaintiffs and not to be used or shared outside of Lancium” and that “[Mr.] McNamara gave his word that he would abide by this confidentiality.” (D.I. 19 at ¶ 71) The claim then asserts that Plaintiffs relied on Mr. McNamara’s assurances that he would “keep the BearBox Technology confidential” and that Plaintiffs would not have continued to share information with

Mr. McNamara had they known that Defendants would violate that promise and incorporate the BearBox Technology into patent applications. (*Id.* at ¶¶ 71-72) As for damages suffered as a result of Defendants’ alleged misrepresentation, Plaintiffs allege the loss of potential patent rights as well as the “*costs of Plaintiffs’ know-how converted under the guise of a potential business relationship.*” (*Id.* at ¶ 73 (emphasis added))

Thus, here the claim as pleaded seems to turn in significant part on whether Mr. McNamara made a representation about confidentiality to Plaintiffs and whether that representation was false; in resolving those questions, whether Mr. Storms was a true inventor is not necessarily a relevant consideration. And Plaintiffs’ request for relief is directed, at least in part, to damages for something other than the loss of the ability to license the patent at issue. In light of this, and in light of the lack of caselaw finding negligent misrepresentation claims to be preempted by federal patent law,¹⁰ the Court concludes that Plaintiffs’ negligent misrepresentation claim is not preempted. *Cf. HIF Bio, Inc. v. Yung Shin Pharms. Indus. Co.*, 600 F.3d 1347, 1356 (Fed. Cir. 2010) (finding that plaintiffs’ causes of action for, *inter alia*, fraud and intentional and negligent interference with contractual relations claims were not preempted, where the inventorship issue was not essential to the resolution of such claims).

However, the Court agrees with Defendants that Count V should nevertheless be dismissed. In this regard, Defendants rightly contend that the claim is fatally flawed because Plaintiffs fail to plead a legal duty on the part of Lancium/Mr. McNamara. (D.I. 33 at 18-20; D.I. 46 at 10)

¹⁰ Indeed, Defendants acknowledge that they are not aware of any case that specifically addresses preemption of a claim for negligent misrepresentation by federal patent law. (D.I. 33 at 13 & n.2)

Whether these Defendants owed a duty to Plaintiffs is a question of law. *Barrie v. V.P. Exterminators, Inc.*, 625 So. 2d 1007, 1015 (La. 1993). “[T]he initial inquiry is whether, as a matter of law, a duty is owed to this particular plaintiff to protect him from this particular harm[.]” and “[a] negative answer . . . results in a determination of no liability.” *Miller v. Lowe*, Civil Action No. 08-1624, 2009 WL 4730201, at *4 (W.D. La. Dec. 4, 2009) (internal quotation marks and citations omitted).

Plaintiffs, for their part, argue that the Amended Complaint sufficiently pleads “the existence of a duty stemming from the confidential relationship established between the parties.” (D.I. 42 at 15) In this regard, the Supreme Court of Louisiana has explained that a confidential or fiduciary relationship giving rise to the requisite duty can exist between “generally all persons who are associated by any relation of trust and confidence” such as, for example, “trustee and beneficiary, attorney and client, parent and child, or husband and wife . . . [and] partners and co-partners, principal and agent, master and servant, physician and patient[.]” *Bunge Corp. v. GATX Corp.*, 557 So. 2d 1376, 1384 n.4 (La. 1990).

None of these types of relationships (or something close to them) are at issue here. Instead, the Amended Complaint pleads that Mr. Storms and Mr. McNamara were simply two men with no previous relationship who met each other at a conference, had dinner and went on to discuss a “potential business relationship”—and that in those conversations, Mr. Storms shared what he considered to be confidential information with Mr. McNamara. (D.I. 19 at ¶¶ 32, 72-73) Plaintiffs cite to no cases suggesting that this type of “business relationship” gives rise to the requisite duty under Louisiana law. And courts applying Louisiana law to circumstances involving similar types of business relationships have found the opposite. *See, e.g., House of Raeford Farms of La. LLC v. Poole*, CIVIL ACTION NO. 19-271, 2021 WL 1081837, at *4-6

(W.D. La. Mar. 18, 2021) (dismissing a fraudulent misrepresentation claim, which similarly requires that a special relationship exist between a plaintiff and defendant establishing a duty to disclose omitted information, where the parties were in a seller-buyer relationship, since fiduciary duties do not arise from “ordinary supplier-customer contracts”) (internal quotation marks and citation omitted); *S. Serv. Corp. v. Tidy Bldg. Servs., Inc.*, No. Civ.A. 04-1362, 2004 WL 2784909, at *6 (E.D. La. Dec. 1, 2004) (finding that the plaintiff’s complaint failed to state a claim for negligent misrepresentation, where “[t]he relationships between Tidy and its customers and Tidy and [plaintiff] do not fall within the class of special relationships of trust and confidence; they are ordinary commercial relationships”). Plaintiffs’ negligent misrepresentation claim therefore fails as a matter of law.

D. Conclusion

In sum, the Court agrees with Defendants that Plaintiffs’ conversion and unjust enrichment claims are preempted by federal patent law, and that Plaintiffs’ negligent misrepresentation claim is wanting as a matter of law due to the failure to sufficiently allege the requisite duty. Plaintiffs have requested leave to amend the Amended Complaint if Defendants’ Motion is granted. (D.I. 42 at 18) The Court believes that this request should be granted as to Counts III and IV because: (1) it is not clear that allowing the opportunity to amend would be a futile act;¹¹ (2) this is the first time the Court has found Plaintiffs’ claims to be deficiently pleaded; and (3) leave to amend should be given freely “when justice so requires.” Fed. R. Civ. P. 15(a)(2). As to Count V, however, the Court does not see how Plaintiffs can plead a plausible

¹¹ If Plaintiffs are given the opportunity to replead, they should be mindful of Defendants’ other arguments for dismissal of these Counts, (D.I. 33 at 14-18), and should address those arguments in the new pleading to the extent they need to, as it is not likely that Plaintiffs would be given a further chance to replead thereafter.

claim (i.e. how Plaintiffs can possibly allege the requisite duty). Thus, the Court also recommends that dismissal of Counts III and IV be without prejudice and that dismissal of Count V be with prejudice. The Court suggests that, to the extent the District Court affirms the Court's recommendation, Plaintiffs be given leave to file one further amended complaint within 14 days. *TriDiNetworks Ltd. v. Signify N. Am. Corp.*, Civil Action No. 19-1063-CFC-CJB, 2020 WL 2839224, at *5 (D. Del. June 1, 2020).

IV. CONCLUSION

For the foregoing reasons, the Court recommends that Defendants' Motion be GRANTED.¹²

This Report and Recommendation is filed pursuant to 28 U.S.C. § 636(b)(1)(B), Fed. R. Civ. P. 72(b)(1), and D. Del. LR 72.1. The parties may serve and file specific written objections within fourteen (14) days after being served with a copy of this Report and Recommendation. Fed. R. Civ. P. 72(b)(2). The failure of a party to object to legal conclusions may result in the loss of the right to *de novo* review in the district court. *See Sincavage v. Barnhart*, 171 F. App'x 924, 925 n.1 (3d Cir. 2006); *Henderson v. Carlson*, 812 F.2d 874, 878-79 (3d Cir. 1987).

The parties are directed to the Court's Standing Order for Objections Filed Under Fed. R. Civ. P. 72, dated October 9, 2013, a copy of which is available on the District Court's website, located at <http://www.ded.uscourts.gov>.

Dated: January 18, 2022


Christopher J. Burke
UNITED STATES MAGISTRATE JUDGE

¹² Defendants' request for oral argument, (D.I. 47), is DENIED.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEAR BOX LLC and AUSTIN STORMS,)
)
 Plaintiffs,)
)
 v.) C.A. No. 21-534(MN) CJB)
)
 LANCIUM LLC, MICHAEL T.)
 MCNAMARA and RAYMOND E. CLINE,)
 JR.,)
)
 Defendants.)

ORDER ADOPTING REPORT AND RECOMMENDATION


At Wilmington, this 2nd day of February 2022:

WHEREAS, on January 18, 2022, Magistrate Judge Burke issued a Report and Recommendation (“the Report”) (D.I. 92) in this action recommending that the Court grant Defendants’ Motion for Judgment on the Pleadings (D.I. 32); and

WHEREAS, no party filed objections to the Report pursuant to Rule 72(b)(2) of the Federal Rules of Civil Procedure in the prescribed period, and the Court finding no clear error on the face of the record.

THEREFORE, IT IS HEREBY ORDERED, that the Report and Recommendation (D.I. 92) is ADOPTED and Defendants' Motion for Judgment on the Pleadings (D.I. 32) is GRANTED. Counts III and IV of Plaintiff's Amended Complaint (D.I. 19) are DISMISSED WITHOUT PREJUDICE and Count V is DISMISSED WITH PREJUDICE.

Plaintiff is given leave to file a further amended complaint on or before February 16, 2022.


The Honorable Maryellen Noreika
United States District Court Judge

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,)	
)	
Plaintiffs,)	
)	
v.)	C.A. No. 21-534-MN-CJB
)	
LANCIUM LLC, MICHAEL T.)	JURY TRIAL DEMANDED
MCNAMARA, and RAYMOND E. CLINE,)	
JR.)	
)	
Defendants.)	

SECOND AMENDED COMPLAINT

Plaintiffs BearBox LLC (“BearBox”) and Austin Storms (collectively, “Plaintiffs” or “BearBox”) bring this action against Lancium LLC (“Lancium”), Michael T. McNamara, and Raymond E. Cline, Jr. (collectively “Defendants”) to correct the inventorship of U.S. Patent No. 10,608,433 (the “’433 Patent”) and to recover damages, injunctive relief, declaratory relief, and other remedies for Defendants’ wrongful actions to obtain, convert, misuse, disclose, and claim as their own Plaintiffs’ proprietary cryptocurrency mining technology. Plaintiffs further allege as follows:

INTRODUCTION

1. This case is about the Defendants’ theft and unauthorized use of system designs, data, know-how, and physical documents describing those designs, data, and know-how, as well as inventions that rightfully belong to Plaintiffs.

2. Plaintiffs developed proprietary technology relating to cryptocurrency mining systems. By way of background, BearBox’s technology generally relates to an energy-efficient cryptocurrency mining system and related methods that reduce the inefficiency and

environmental impact of energy-intensive mining operations by better utilizing available energy resources to increase the stability of the energy grid, minimize a mining operation's impact on peak-demand, and also alleviate energy oversupply and undersupply conditions. BearBox's technology can be used to mine cryptocurrency, such as Bitcoin.

3. The Defendants deceptively induced the Plaintiffs to disclose BearBox’s technology to them under the guise of a possible business deal or “collaboration,” a word used by McNamara in encouraging Plaintiffs’ disclosures, between Defendants and Plaintiffs to jointly commercialize the BearBox Technology.

4. The Defendants stole BearBox's technology, including system designs, documents, data, and know-how, from Plaintiffs by converting it and using it to modify Defendants' software, e.g. Smart Response™, to operate as BearBox's technology did, and then selling, licensing, seeking investments in support of, and otherwise monetizing that modified Smart Response™ software for great profit.

5. Defendants went even further, by filing a U.S. patent application, which later matured into the '433 Patent, that wrongfully disclosed portions of BearBox's technology to the U.S. Patent and Trademark Office and ultimately to the public as technology that was invented by Defendants. The claimed subject matter of the '433 Patent falls fully within the scope of BearBox's technology, though is not a complete representation of the BearBox technology. By obtaining the '433 Patent with claims directed to portions of BearBox's technology, the Defendants have wrongfully obtained a patent that prevents Plaintiffs from building their own system, or otherwise monetizing their system designs, data, documents, and know-how.

6. Plaintiffs bring this action to recover damages caused by Defendants' theft and unauthorized use, and subsequent exploitation of Plaintiffs system design, data, documents, and know-how.

7. Plaintiffs also bring this action to correct the named inventors on the '433 Patent. The inventions claimed in the '433 Patent are inventions conceived by Storms, founder and president of BearBox.

PARTIES

8. BearBox LLC is a limited liability company organized and existing under the laws of Louisiana with its principal place of business at 4422 Highway 22, Mandeville, Louisiana 70471.

9. Plaintiff Austin Storms is an individual residing in Mandeville, Louisiana.

10. On information and belief, Defendant Lancium is a Delaware limited liability company with its principal place of business at 6006 Thomas Rd, Houston, Texas 77041. On information and belief, Lancium has a registered agent capable of accepting service in this district, Harvard Business Services, Inc. with a place of business at 16192 Coastal Highway, Lewes, DE 19958.

11. On information and belief, Defendant Michael T. McNamara is the Chief Executive Officer and a founder of Lancium and resides in Newport Beach, California. Defendant McNamara is named as a purported inventor on the face of the '433 Patent.

12. On information and belief, Defendant Raymond E. Cline, Jr. is the Chief Computing Officer of Lancium and resides in Houston, Texas. Defendant Cline is named as a purported inventor on the face of the '433 Patent.

JURISDICTION

13. This is an action seeking correction of the named inventors of a United States patent under 35 U.S.C. § 256. As such, this action arises under the laws of the United States.

14. This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under an Act of Congress relating to patents, specifically 35 U.S.C. § 256.

15. The Court has supplemental jurisdiction under 28 U.S.C. § 1367 over all asserted claims under state law because those claims are so related to the claims in this action that arise under federal law that they form part of the same case or controversy.

16. The Court also has jurisdiction pursuant to 28 U.S.C. § 1332, as complete diversity of citizenship exists among the parties, and the amount in controversy exceeds \$75,000. Plaintiff BearBox is a citizen of the State of Louisiana because it is organized under the laws of the State of Louisiana and has its principal place of business in the State of Louisiana. Plaintiff Storms is a citizen of the State of Louisiana because he resides in the State of Louisiana. In contrast, none of the Defendants are citizens of the State of Louisiana. Defendant Lancium is a citizen of the States of Delaware and Texas because it is organized under the laws of the State of Delaware and has its principal place of business in the State of Texas. Defendant McNamara is a citizen of the State of California because he resides in the State of California. Defendant Cline is a citizen of the State of Texas because he resides in the State of Texas. Therefore, because the Plaintiffs are both citizens of the State of Louisiana (and no other states) for purposes of diversity jurisdiction, and none of the Defendants are citizens of the State of Louisiana, complete diversity exists among the parties.

17. This Court has general personal jurisdiction over Lancium because it is organized under the laws of the State of Delaware and because it maintains an ongoing presence in this District at least through its registered agent.

18. This Court has specific personal jurisdiction over each of Defendants McNamara and Cline at least under Title 6 of the Delaware Code, § 18-109(a).

19. On information and belief, Defendant McNamara is the Chief Executive Officer of Lancium. On information and belief, as the Chief Executive Offer, McNamara participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

20. McNamara is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from the interests of Lancium and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claims against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant McNamara together. Plaintiffs' claims against Defendant McNamara arise out of his exercise of his powers as Chief Executive Officer of Lancium.

21. On information and belief, Defendant Cline is the Chief Computing Officer of Lancium. On information and belief, as the Chief Computing Officer, Cline participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

22. Cline is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from Lancium's interest and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claim against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant Cline together.

Plaintiffs' claims against Defendant Cline arise out of his exercise of his powers as Chief Computing Officer of Lancium.

23. The actions of Defendants McNamara and Cline establish sufficient minimum contacts with Delaware under Delaware law and the United States Constitution to give this Court personal jurisdiction over each of them.

24. As described below, each Defendant has committed acts giving rise to this action.

VENUE

25. Venue is proper in this District under 28 U.S.C. § 1391(b)(3) because there is no district in which an action may otherwise be brought as provided in § 1391(b) and Defendant Lancium is subject to the Court's personal jurisdiction with respect to this action.

PLAINTIFFS' PROPRIETARY CRYPTOCURRENCY MINING TECHNOLOGY

26. As of 2018, the amount of energy required to process computer algorithms to mine cryptocurrencies like Bitcoin was three times greater than the energy required to physically mine gold. Conventional mining of "copper, gold, platinum, and rare earth oxides are 4, 5, 7, and 9 megajoules to generate one U.S. dollars," while "it costs an average of 17 megajoules to mine \$1 worth of bitcoin."¹ The large amount of energy required to mine cryptocurrencies can make such mining financially prohibitive, and even when financially lucrative, the energy required for cryptocurrency mining can be harmful to the environment, with studies showing potential carbon dioxide emissions from cryptocurrency mining "single-handedly rais[ing] global temperatures by 2 degrees by 2023." *Id.*

¹ <https://www.marketwatch.com/story/mining-bitcoin-is-3-times-more-expensive-than-mining-gold-research-paper-finds-2018-11-06>

27. At the same time, some forms of electrical power generation, such as wind, solar, or other types of variable renewable energy sources, are inefficient and suffer from a timing imbalance between peak demand and energy production. Additionally, these non-dispatchable renewable energy sources cannot be controlled by operators in response to dynamically changing grid conditions. As a result, these renewable energy sources and their operators frequently sell power at low or negative prices, sometimes even being curtailed or shut down entirely, until demand increases and grid conditions change.

28. Because cryptocurrency mining using proof-of-work is a computationally demanding process, it requires a significant amount of energy to operate. As a result, industrial-scale cryptocurrency mining places a large energy burden on the power grid, driving demand and costs as well as increasing the likelihood of transmission or other grid component failure if not managed correctly.

29. In late 2018 and early 2019, Austin Storms sought to address these problems by developing energy-efficient cryptocurrency mining systems and methods that reduce the environmental impact of energy-intensive mining operations. Storms conceived of a system that better uses available energy resources to increase the stability of the energy grid, minimize a mining operation's impact on peak-demand, and help alleviate energy oversupply and undersupply conditions, all while decreasing the overall energy costs of the mining operation and increasing its profitability.

30. Austin Storms conceived of and developed BearBox's technology. Storms is the president and founder of BearBox. BearBox's technology includes hardware and software components. Structurally, BearBox's technology includes a housing for a plurality of miners (such as ASICs, graphics cards, or the like) under the direction of a smart controller(s).

31. In BearBox's technology, a controller (such as a power distribution unit, network interface, computer, and/or the like) monitors various external factors, such as current and expected energy demand/pricing information, current and expected cryptocurrency pricing, current and expected network hashrate and difficulty, and the like. Based on these external factors, the controller(s) determines appropriate times to mine cryptocurrency in accordance with a desired performance strategy (for example, profitability thresholds). At the appropriate times, the controller initiates mining, for example, by powering on a number of miners based on various conditions internal and external conditions.

DEFENDANTS WRONGFULLY STOLE BEARBOX'S TECHNOLOGY

32. In May 2019, Storms attended the Fidelity FCAT Mining Summit in Boston, Massachusetts on behalf of BearBox to promote BearBox's technology and seek potential customers for his revolutionary system.

33. While at the conference, Storms met Defendant McNamara. Defendant McNamara showed immediate interest in BearBox's technology. Under the rouse of a potential business relationship, McNamara deceptively pumped Storms for details about BearBox's technology over the course of several exchanges, which included conversations, emails, and text messages about BearBox's technology. Storms took McNamara to dinner where McNamara continued to pump Storms for details about BearBox's technology.

34. Following the conference, McNamara continued to press Storms for additional details about BearBox's technology via text messaging and email. Storms described BearBox's technology and provided annotated system diagrams, component specifications, and modeled data sets to mimic real-world Bitcoin variables such as Bitcoin price and network hashrate, energy prices, time intervals, and power thresholds, and computed profitability figures. Storms

communications to McNamara about BearBox's technology was for the sole purpose of potential partnership. Storms included express confidentiality notices in his email communications with Defendant McNamara.

35. After Storms disclosed BearBox's technology to McNamara, McNamara abruptly ended all communications with Storms.

36. Storms last communicated with McNamara on May 9, 2019 via e-mail, and after sending that message, Storms did not hear from McNamara again.

37. At that time, Storms understood that McNamara was not interested in investing in or collaborating with BearBox's technology. But Storms had no reason to suspect that Defendants were being deceptive and would steal BearBox's technology, including the system designs, diagrams, data, documents, and know-how, and use it to build Defendants' own system that would function as BearBox's technology did.

38. In the days, weeks, and months following Plaintiffs' disclosures to McNamara, Defendants began working with partners, outside consultants, and vendors to help Defendants bring BearBox's technology to market, labeled as Defendants' own system. Defendants efforts included working internally to modify their Smart Response™ software to work as BearBox's technology did.

39. Defendants soon became confident they could exploit BearBox's technology for profit by modifying their Smart Response™ software, and then using, selling, licensing, and otherwise monetizing that modified Smart Response™ software, but Defendants also realized that others would eventually learn of Defendants' new-and-improved Smart Response™ software, now with the benefits of BearBox's technology, and may copy it in order to compete with Defendants.

40. Defendants believed that the profits they would realize from exploiting their modified Smart Response™ software would be significantly reduced if outside competitors were allowed to copy Defendants' unauthorized use of BearBox's technology and compete directly with Defendants.

41. Because of these realizations, Defendants hatched a plan to patent some of the profitable attributes of BearBox's technology. If such a patent was obtained, Defendants believed that such a patent would effectively grant Defendants a monopoly to exploit significant portions of BearBox's technology for profits that would far exceed any profits should Defendants be left to compete with others, including BearBox.

42. On information and belief, Defendants filed U.S. provisional patent application No. 62/927,119 on October 28, 2019, naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

43. In addition to falsely claiming to be the inventors of the inventions disclosed in the application, Defendants wrongfully disclosed, without authorization, portions of BearBox's technology to the United States Patent and Trademark Office.

44. Likewise, on December 4, 2019, Defendants filed U.S. Patent Application Serial No. 16/702,931, once again naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

45. The '433 Patent issued on March 31, 2020 naming Defendants McNamara and Cline as the sole purported inventors on the face of the patent. A true and correct copy of the '433 Patent is attached hereto as Exhibit A.

46. The inventions claimed in the '433 patent fall within the scope of BearBox's technology, yet Defendants falsely identified themselves as the inventors of the claimed

inventions, when, in fact, Storms is the sole inventor of the claimed inventions. Not all aspects of BearBox's technology that was stolen and used by Defendants was described and claimed in the '433 Patent. For example, Defendants theft and unauthorized use included confidential aspects of BearBox technology's system design, diagrams, data, and know-how to modify Defendants' Smart Response software to develop and exploit, among other things, methods for energy value arbitrage, including cryptocurrency mining systems and arbitraging related energy values.

47. On information and belief, McNamara and Cline assigned their purported rights in the '433 patent to Lancium. On information and belief, at all times, Lancium was aware that McNamara and Cline, both officers of Lancium, were not the rightful inventors of BearBox's technology disclosed in the patent and the inventions claimed in the patent.

48. Defendants McNamara and Cline each submitted signed declarations falsely swearing that they were "an original joint inventor" of the claimed subject matter. A true and correct copy of Defendant McNamara's and Defendant Cline's declarations are attached as Exhibit B.

49. On June 19, 2020, Defendants announced that its Smart Response™ software allowed it to qualify as the "first successful load-only Controllable Load Resource (CLR) designation," which Defendants described as a "breakthrough achievement," that allows for dispatching "excess energy into the grid during times when energy is most expensive," adding that "with this revolutionary advancement, data centers will also be able to earn additional revenue providing ancillary services to ERCOT."

50. Since then, Defendants have used BearBox's technology, including the stolen system designs, diagrams, data, and know-how, and its subsequently-modified Smart Response™ software to function as it did, to allow Defendants to profit significantly from use,

license, sale, and investments related to its modified Smart Response™ software. Defendants are not, and have never, been entitled to these profits.

51. In addition to illegally profiting from its theft and subsequent use of BearBox's technology, Defendants made use of its ill-gotten '433 patent.

52. On August 14, 2020, Lancium filed a lawsuit in the U.S. District Court for the Western District of Texas against Layer1 Technologies, Inc. ("Layer1") asserting that Layer1 infringes the '433 patent, and stating that "Lancium's Controllable Load Resource Technology attracted the interest of many companies, including, upon information and belief, Layer1." That case is captioned *Lancium LLC v. Layer1 Technologies, Inc.*, Case No. 6:20-cv-739 (W.D. Texas) (the "Layer1 Lawsuit").

53. As part of the Layer1 Lawsuit, Defendants falsely asserted that McNamara and Cline are the sole inventors of the inventions claimed in the '433 patent.

54. Plaintiffs became aware of Defendants' wrongful use of BearBox's technology on or about August 17, 2020, when they learned about the Layer1 Lawsuit through a press release dated August 14, 2020, posted by Lancium on PRNewswire. That press release is available at the following URL: <https://www.prnewswire.com/news-releases/controllable-load-resource-clr-market-leader-lancium-files-patent-infringement-lawsuit-against-layer1-301112687.html>.

55. Before seeing the August 14, 2020 press release, Plaintiffs were unaware of Defendants' wrongful use of BearBox's technology and was unaware of the '433 patent.

56. On March 5, 2021, Lancium and Layer1 entered a Stipulation to Dismiss with Prejudice in the Layer1 Lawsuit. According to the stipulation, the parties had entered a Settlement Agreement to resolve the Layer1 Lawsuit.

57. According to a press release issued by Lancium on March 8, 2021, Lancium and Layer 1 “have entered into a mutually beneficial partnership. Layer1 has licensed Lancium’s intellectual property and Lancium will provide Smart Response™ software and services to Layer1.” The press release is available at the following URL: <https://www.prnewswire.com/news-releases/lancium-and-layer1-settle-patent-infringement-suit-301242602.html>.

**COUNT I
CORRECTION OF INVENTORSHIP FOR THE ’433 PATENT:
AUSTIN STORMS AS SOLE INVENTOR**

58. Plaintiffs incorporate the above paragraphs by reference.

59. Storms is the sole inventor of the subject matter claimed in the ’433 Patent.

60. Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the ’433 patent and the currently listed inventors on the ’433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms or BearBox.

61. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the ’433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

**COUNT II
IN THE ALTERNATIVE, CORRECTION OF INVENTORSHIP FOR THE ’433
PATENT: AUSTIN STORMS AS JOINT INVENTOR WITH THE CURRENTLY
NAMED INVENTORS**

62. Plaintiffs incorporates the above paragraphs by reference.

63. In the alternative, Storms is a joint inventor of the subject matter claimed in the ’433 Patent and should be added to the individuals currently named as inventors on the ’433 Patent.

64. Through omission, inadvertence, and/or error, Storms was not listed as an

65. Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting

COUNT III

66. Plaintiffs incorporate the above paragraphs by reference.

67. At all relevant times, Plaintiffs and Defendants were engaged in interstate

68. In the course of the interactions between Austin Storms and Defendants,

69. Defendants knew the information they received from Austin Storms about BearBox's technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship, and was in no way the information received to be used by Defendants' for their own profit. McNamara himself referred to the exchanges with Plaintiffs' "lots of stuff to collaborate on." Defendants further knew the technical information they received from Storms about BearBox's technology was confidential based upon Defendants' own actions taken to protect that information once they possessed it, and misappropriated it by using in modification of their Smart Response™ software. Defendants' restricted access to the information, including as it was represented in the Smart Response™ software, required confidentiality agreements for third party customers, potential customers, and licensees, and other internal measures were taken.

70. Plaintiffs included a confidentiality designation on communications to which the confidential trade secret information was attached. Plaintiffs took additional reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of individuals that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its computers. The confidential information was, at all other times, secured on a password protected computer. The confidential information was not externally accessible through the internet. When that password protected computer was not attended, it was secured in a locked building at Storms residence, which included other security measures such as a concierge and elevator keys.

71. For the very limited number of additional occasions the confidential, trade secret information was shared, it was shared with a person or entity subject to contractual obligations regarding confidentiality and/or prohibitions on unauthorized use. For the limited number of occasions BearBox shared the confidential, trade secret information with potential customers, it was done only with BearBox management approval.

72. Defendants misappropriated Plaintiffs' trade secrets when they used BearBox technology, without Plaintiffs' authorization, in at least its Smart Response™ software. Defendants have disclosed and/or used the confidential information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, and using it as Defendants' own, and without permission was improper. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

73. As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages or other financial harm in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, financial loss for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability.

74. The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D).

**COUNT IV
TRADE SECRET MISAPPROPRIATION UNDER
LOUISIANA CIVIL PRACTICE AND REMEDIES CODE, TITLE 51, SECTION 13-A**

75. Plaintiffs incorporate the above paragraphs by reference.

76. At all relevant times, Plaintiffs and Defendants were engaged in trade or commerce within the meaning of La. R.S. 51:1431-39 (2011), the Louisiana Uniform Trade Secrets Act. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods.

77. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about BearBox's technology relating to certain features, such as methods for energy value arbitrage, including cryptocurrency mining systems and arbitrating related energy values. Such information constitutes trade secrets of substantial economic value, at least because methods for energy value arbitrage, including cryptocurrency mining systems and arbitrating related energy values may be used for substantial profit. This confidential technical information was not known to the public or to other persons who can obtain economic value from its disclosure or use. This information constituted a trade secret under the Louisiana Uniform Trade Secret Act.

78. Defendants knew the technical information they received from Austin Storms about BearBox's technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship, and was in no way the information received to be used by Defendants' for their own profit. Defendants further knew the information they received from Storms about BearBox's technology was confidential based upon Defendants' own actions taken to protect that information once they possessed it. Defendants' restricted access to the

information, required confidentiality agreements for third party customers, potential customers, and licensees, and other internal measures were taken. Defendants further took these actions to protect Plaintiffs' confidential information after Defendants misappropriated it by including it for use in Defendants' Smart Response™ software.

79. Plaintiffs included a confidentiality designation on communications to which the confidential trade secret information was attached. Plaintiffs took additional reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its computers. The confidential information was, at all other times, secured on a password protected computer. The confidential information was not externally accessible through the internet. When that password protected computer was not attended, it was secured in a locked building.

80. For the very limited number of additional occasions the confidential, trade secret information was shared, it was shared with a person or entity subject to contractual obligations regarding confidentiality and/or prohibitions on unauthorized use. For the limited number of occasions BearBox shared the confidential, trade secret information with potential customers, it was done only with BearBox management approval.

81. Defendants misappropriated Plaintiffs' trade secrets when they used BearBox technology, without Plaintiffs' authorization, in at least its Smart Response™ software. Defendants have disclosed and/or used the confidential technical information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, and using it as Defendants' own, and without permission was improper. Further, Defendant McNamara

agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret.

Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

82. As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages or other financial harm in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, financial loss for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability.

83. The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under § 51:1434.

COUNT V CONVERSION BY LANCIUM, MCNAMARA, AND CLINE

84. Plaintiffs incorporates the above paragraphs by reference.

85. Austin Storms, in his capacity as founder and President of BearBox, conceived, developed, and reduced to practice BearBox's technology. Plaintiffs own BearBox's technology, including system designs, documents, data, and know-how. Plaintiffs owned this property during all relevant time periods in this suit.

86. Defendants induced Plaintiffs to provide information and documents regarding BearBox's technology to Defendants by misrepresenting that Defendants were interested in investing in or otherwise collaborating with BearBox. Information and documents regarding BearBox's technology were provided to Defendants solely for the purposes of evaluation for a potential business relationship.

87. Without Plaintiffs' consent, Defendants intentionally and willfully assumed dominion and control over BearBox's technology, including system designs, documents, data, and know-how, and improperly used it to modify their Smart Response™ software, and corresponding system designs, to function as reflected in BearBox's system designs, documents, data, and know-how, and subsequently used, sold, licensed, and procured investments related to, and otherwise monetized, that software for substantial profit.

88. Defendants' unlawful exercise of dominion over confidential Bearbox technology, including system designs, documents, data, and know-how not otherwise found to be a trade secret, or having value beyond the value of the trade secret(s), has permanently interfered with Plaintiffs' valuable property rights.

89. Despite providing Defendants with the system designs, documents, data, and know-how that allowed Defendants to modify their Smart Response™ software, and corresponding system designs, Defendants have not compensated or recognized Plaintiffs for the use of BearBox's technology. Defendants' actions constitute an improper and unauthorized use of Plaintiffs' property.

90. As a result of Defendants' improper and unauthorized use of Plaintiffs' system designs, documents, data, and know-how to reconstruct and use Bearbox's technology, Plaintiffs have suffered and will continue to suffer damages and other financial harms in an amount to be proven at trial. As such, Plaintiffs are entitled to damages resulting from Defendants' improper and unauthorized use.

COUNT VI UNJUST ENRICHMENT BY LANCIUM, MCNAMARA, AND CLINE

91. Plaintiffs incorporate the above paragraphs by reference.

92. Plaintiffs conferred a benefit on Defendants by providing them valuable technology, specifically BearBox's technology, including system designs, documents, data, and know-how.

93. Defendants pressed Storms to provide more information and materials, having recognized the benefit that Defendants received by having access to system designs, documents, data, and know-how reflecting BearBox's technology.

94. Defendants induced Plaintiffs to provide these system designs, documents, data, and know-how by misrepresenting that Defendants were interested in investing in or otherwise collaborating with Bearbox. Defendants' actions in obtaining Plaintiffs' system designs, documents, data, and know-how were deceptive.

95. Defendants retained BearBox's technology, including system designs, documents, data, and know-how, and used it to their own advantage, and at BearBox's expense.

96. Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct. For example, without Plaintiffs' consent, Defendants used Plaintiffs' system designs, documents, data, and know-how to modify their Smart Response™ software to function as BearBox's technology did. As a result, Defendants are deriving an unjust benefit from exploiting BearBox's property.

97. Despite providing Defendants with the system designs, documents, data, and know-how that allowed Defendants to modify their system designs and Smart Response™ software to function as BearBox's technology did, Defendants have not compensated Plaintiffs for the use of BearBox's technology. It would be inequitable for Defendants to retain these benefits under these circumstances.

98. Defendants' unauthorized use of Bearbox's technology has impoverished Plaintiffs by depriving them of the fruits of their labor. In addition to the costs incurred and labor expended in creating Bearbox's technology, Plaintiffs are no longer able to obtain any economic benefit from any use of Bearbox's technology.

99. Defendants' enrichment through the use of Plaintiffs' system designs, documents, data, and know-how to reconstruct and use Bearbox's technology is directly connected to Plaintiffs' impoverishment because Defendants have failed to compensate Plaintiffs for their efforts.

100. There is no justification at law or in contract for Defendants' uncompensated use of Plaintiffs' unique and novel invention.

101. Because Plaintiffs have incurred, and continue to incur, detriment in the form of loss of money and property as a result of Defendants' wrongful use of system designs, documents, data, and know-how reflecting BearBox's technology, Plaintiffs are entitled to compensation to the extent Defendants have been enriched or Plaintiffs have been impoverished.

JURY DEMAND

102. Under Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiffs respectfully demand a trial by jury on all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, BearBox respectfully requests the following relief:

A. An order that the Director of the United States Patent and Trademark Office correct the inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

Patent;

joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

violation of the United States federal patent laws;

based on, and/or derived in whole or in part from the use of BearBox's technology;

BearBox's technology;

appropriate financial relief, all in an amount to be determined at trial, with interest;

their actions set forth in this Complaint;

I. A finding that this is an exceptional case warranting imposition of attorney fees against Defendants and an award to Plaintiffs of its reasonable costs and attorney fees incurred in bringing this action pursuant to 35 U.S.C. § 285;

J. An award of exemplary damages and/or an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D) and/or under La. R.S. § 51:1434; and

K. An award of such further relief at law or in equity, such as preliminary and/or permanent injunctive relief, as the Court deems just and proper.

ASHBY & GEDDES

/s/ Andrew C. Mayo

Andrew C. Mayo (#5207)
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19899
(302) 654-1888
amayo@ashbygeddes.com

*Attorneys for Plaintiffs
BearBox LLC and Austin Storms*

Of Counsel:

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati, III
Chelsea M. Murray
MARSHALL, GERSTEIN & BORUN LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
(312) 474-6300

Dated: February 16, 2022

EXHIBIT A



US010608433B1

(12) **United States Patent**
McNamara et al.

(10) **Patent No.:** **US 10,608,433 B1**
(45) **Date of Patent:** **Mar. 31, 2020**

(54) **METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT**

7,143,300 B2 11/2006 Potter et al.
7,647,516 B2 1/2010 Ranganathan et al.
(Continued)

FOREIGN PATENT DOCUMENTS

(71) Applicant: **Lancium LLC**, Houston, TX (US)

CN 103163904 A 6/2013
KR 20090012523 A 2/2009
WO 2015199629 A1 12/2015

(72) Inventors: **Michael T. McNamara**, Newport Beach, CA (US); **Raymond E. Cline, Jr.**, Houston, TX (US)

OTHER PUBLICATIONS

(73) Assignee: **Lancium LLC**, Houston, TX (US)

(*) Notice: Subject to any disclaimer, the term of this patent is extended or adjusted under 35 U.S.C. 154(b) by 0 days.

Bird et al., "Wind and Solar Energy Curtailment: Experience and Practices in the United States," National Renewable Energy Lab (NREL), Technical Report NREL/TP-6A20-60983, Mar. 2014, 58 pages.

(Continued)

(21) Appl. No.: **16/702,931**

(22) Filed: **Dec. 4, 2019**

Primary Examiner — Christopher E. Everett

(74) *Attorney, Agent, or Firm* — McDonnell Boehnen Hulbert & Berghoff LLP

Related U.S. Application Data

(60) Provisional application No. 62/927,119, filed on Oct. 28, 2019.

(51) **Int. Cl.**

H02J 3/14 (2006.01)

H02J 3/00 (2006.01)

G06F 1/3203 (2019.01)

(52) **U.S. Cl.**

CPC **H02J 3/14** (2013.01); **G06F 1/3203** (2013.01); **H02J 3/008** (2013.01)

(58) **Field of Classification Search**

CPC H02J 3/14; H02J 3/008; G06F 1/3203

See application file for complete search history.

(57)

ABSTRACT

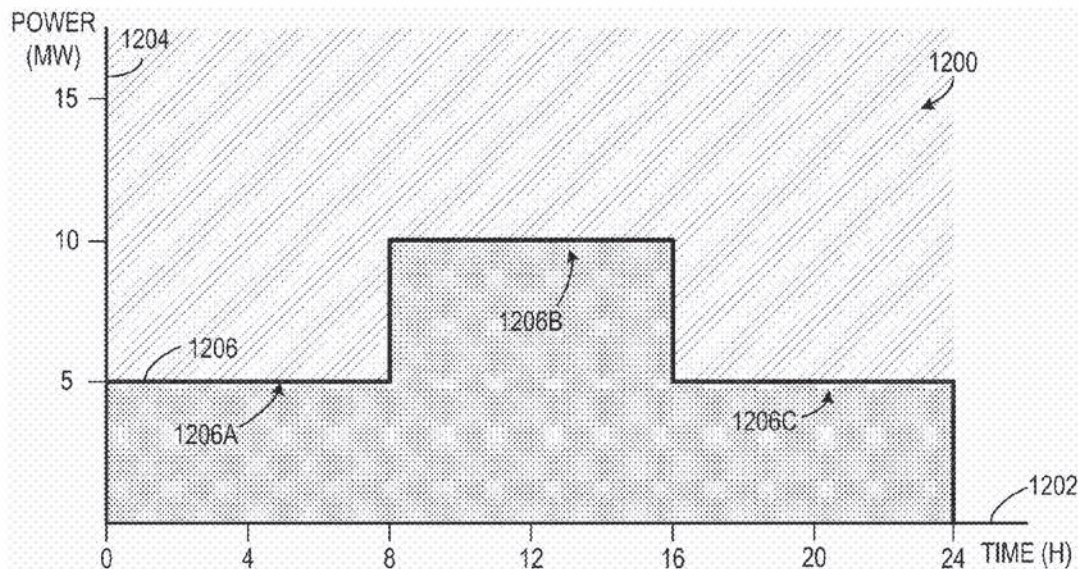
Examples relate to adjusting load power consumption based on a power option agreement. A computing system may receive power option data that is based on a power option agreement and specify minimum power thresholds associated with time intervals. The computing system may determine a performance strategy for a load (e.g., set of computing systems) based on a combination of the power option data and one or more monitored conditions. The performance strategy may specify a power consumption target for the load for each time interval such that each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The computing system may provide instructions the set of computing systems to perform one or more computational operations based on the performance strategy.

(56) **References Cited**

U.S. PATENT DOCUMENTS

6,288,456 B1 9/2001 Cratty
6,633,823 B2 10/2003 Bartone et al.

20 Claims, 16 Drawing Sheets



US 10,608,433 B1

Page 2

(56)

References Cited

U.S. PATENT DOCUMENTS

7,702,931	B2	4/2010	Goodrum et al.	
7,779,276	B2	8/2010	Bolan et al.	
7,861,102	B1	12/2010	Ranganathan et al.	
7,921,315	B2	4/2011	Langgood et al.	
7,970,561	B2	6/2011	Pfeiffer	
8,001,403	B2	8/2011	Hamilton et al.	
8,006,108	B2	8/2011	Brey et al.	
8,214,843	B2	7/2012	Boss et al.	
8,374,928	B2	2/2013	Gopisetty et al.	
8,447,993	B2	5/2013	Greene et al.	
8,571,820	B2	10/2013	Pfeiffer	
8,627,123	B2	1/2014	Jain et al.	
8,700,929	B1*	4/2014	Weber	G06F 30/13 713/310
8,789,061	B2	7/2014	Pavel et al.	
8,799,690	B2	8/2014	Dawson et al.	
9,003,211	B2	4/2015	Pfeiffer	
9,003,216	B2	4/2015	Sankar et al.	
9,026,814	B2	5/2015	Aasheim et al.	
9,207,993	B2	12/2015	Jain	
9,218,035	B2	12/2015	Li et al.	
9,552,234	B2	1/2017	Boldyrev et al.	
10,367,353	B1	7/2019	McNamara et al.	
10,367,535	B2	7/2019	Corse et al.	
10,444,818	B1	10/2019	McNamara et al.	
10,452,127	B1	10/2019	McNamara et al.	
10,497,072	B2	12/2019	Hooshmand et al.	
2002/0072868	A1	6/2002	Bartone et al.	
2003/0074464	A1	4/2003	Bohrer et al.	
2005/0203761	A1*	9/2005	Barr	G06F 1/26 713/320
2006/0161765	A1	7/2006	Cromer et al.	
2008/0030078	A1	2/2008	Whitted et al.	
2008/0094797	A1	4/2008	Coglitore et al.	
2009/0055665	A1	2/2009	Maglione et al.	
2009/0070611	A1*	3/2009	Bower, III	G06F 1/3203 713/322
2009/0089595	A1*	4/2009	Brey	G06F 1/3203 713/300
2010/0211810	A1	8/2010	Zacho	
2010/0328849	A1	12/2010	Ewing et al.	
2011/0072289	A1	3/2011	Kato	
2012/0000121	A1	1/2012	Swann	
2012/0072745	A1	3/2012	Ahluwalia et al.	
2012/0300524	A1	11/2012	Fornage et al.	
2012/0324259	A1	12/2012	Aasheim et al.	
2013/0006401	A1	1/2013	Shan	
2013/0063991	A1	3/2013	Xiao et al.	
2013/0086404	A1*	4/2013	Sankar	G06F 1/305 713/324
2013/0187464	A1	7/2013	Smith et al.	
2013/0227139	A1	8/2013	Suffling	
2013/0306276	A1	11/2013	Duchesneau	
2014/0137468	A1	5/2014	Ching	
2014/0379156	A1	12/2014	Kamel et al.	
2015/0121113	A1	4/2015	Ramamurthy et al.	
2015/0155712	A1	6/2015	Mondal	
2015/0229227	A1	8/2015	Aeloiza et al.	
2015/0277410	A1	10/2015	Gupta et al.	
2015/0278968	A1	10/2015	Steven et al.	
2016/0170469	A1	6/2016	Schgal et al.	
2016/0172900	A1	6/2016	Welch, Jr.	
2016/0187906	A1*	6/2016	Bodas	G05B 15/02 700/287

2016/0198656	A1	7/2016	McNamara et al.
2016/0212954	A1	7/2016	Argento
2016/0324077	A1	11/2016	Frantzen et al.
2017/0023969	A1	1/2017	Shows et al.
2017/0104336	A1	4/2017	Elbsat et al.
2017/0261949	A1	9/2017	Hoffmann et al.
2018/0144414	A1	5/2018	Lee et al.
2018/0202825	A1	7/2018	You et al.
2018/0240112	A1	8/2018	Castinado et al.
2018/0366978	A1	12/2018	Matan et al.
2018/0367320	A1	12/2018	Montalvo
2019/0052094	A1	2/2019	Pmsvsv et al.
2019/0168630	A1	6/2019	Mrlik et al.
2019/0258307	A1	8/2019	Shaikh et al.
2019/0280521	A1	9/2019	Lundstrom et al.
2019/0318327	A1	10/2019	Sowell et al.
2019/0324820	A1	10/2019	Krishnan et al.

OTHER PUBLICATIONS

EPEX Spot, "How They Occur, What They Mean," https://www.epexspot.com/en/company-info/basics_of_the_power_market/negative_prices, 2018, 2 pages.

Final Office Action dated Oct. 1, 2019 for U.S. Appl. No. 16/175,246, filed Oct. 30, 2018, 18 pages.

Ghamkhari et al., "Optimal Integration of Renewable Energy Resources in Data Centers with Behind-the-Meter Renewable Generator," Department of Electrical and Computer Engineering Texas Tech University, 2012, pp. 3340-3444.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017955, dated Apr. 30, 2018, 22 pages.

International Search Report and Written Opinion of PCT Application No. PCT/US2018/017950, dated May 31, 2018, 15 pages.

Non-Final Office Action dated Dec. 5, 2019 for U.S. Appl. No. 16/529,360, filed Aug. 1, 2019, 72 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/596,190, filed Oct. 8, 2019, 72 pages.

Non-Final Office Action dated Nov. 14, 2019 for U.S. Appl. No. 16/132,098, filed Sep. 14, 2018, 25 pages.

Non-Final Office Action dated Nov. 21, 2019 for U.S. Appl. No. 16/529,402, filed Aug. 1, 2019, 57 pages.

Non-Final Office Action dated Dec. 11, 2019 on for U.S. Appl. No. 16/132,062, filed Sep. 14, 2018, 17 pages.

Non-Final Office Action dated Dec. 10, 2019 for U.S. Appl. No. 16/528,348, filed Oct. 8, 2019, 33 pages.

Notice of Allowance dated Apr. 2, 2019, for U.S. Appl. No. 16/175,335, filed Oct. 30, 2018, 12 pages.

Notice of Allowance dated Aug. 15, 2019, for U.S. Appl. No. 16/175,146, filed Oct. 30, 2018, 17 pages.

Notice of Allowance dated Jul. 29, 2019, for U.S. Appl. No. 16/245,532, filed Jan. 11, 2019, 13 pages.

Rahimi, Farrokh, "Using a Transactive Energy Framework," IEEE Electrification Magazine, Dec. 2016, pp. 23-29.

Soluna, "Powering the Block Chain," Aug. 2018, version 1.1, 29 pages.

Wilson, Joseph Nathanael, "A Utility-Scale Deployment Project of Behind-the-Meter Energy Storage for Use in Ancillary Services, Energy Resiliency, Grid Infrastructure Investment Deferral, and Demand-Response Integration," Portland State University, 2016, 154 pages.

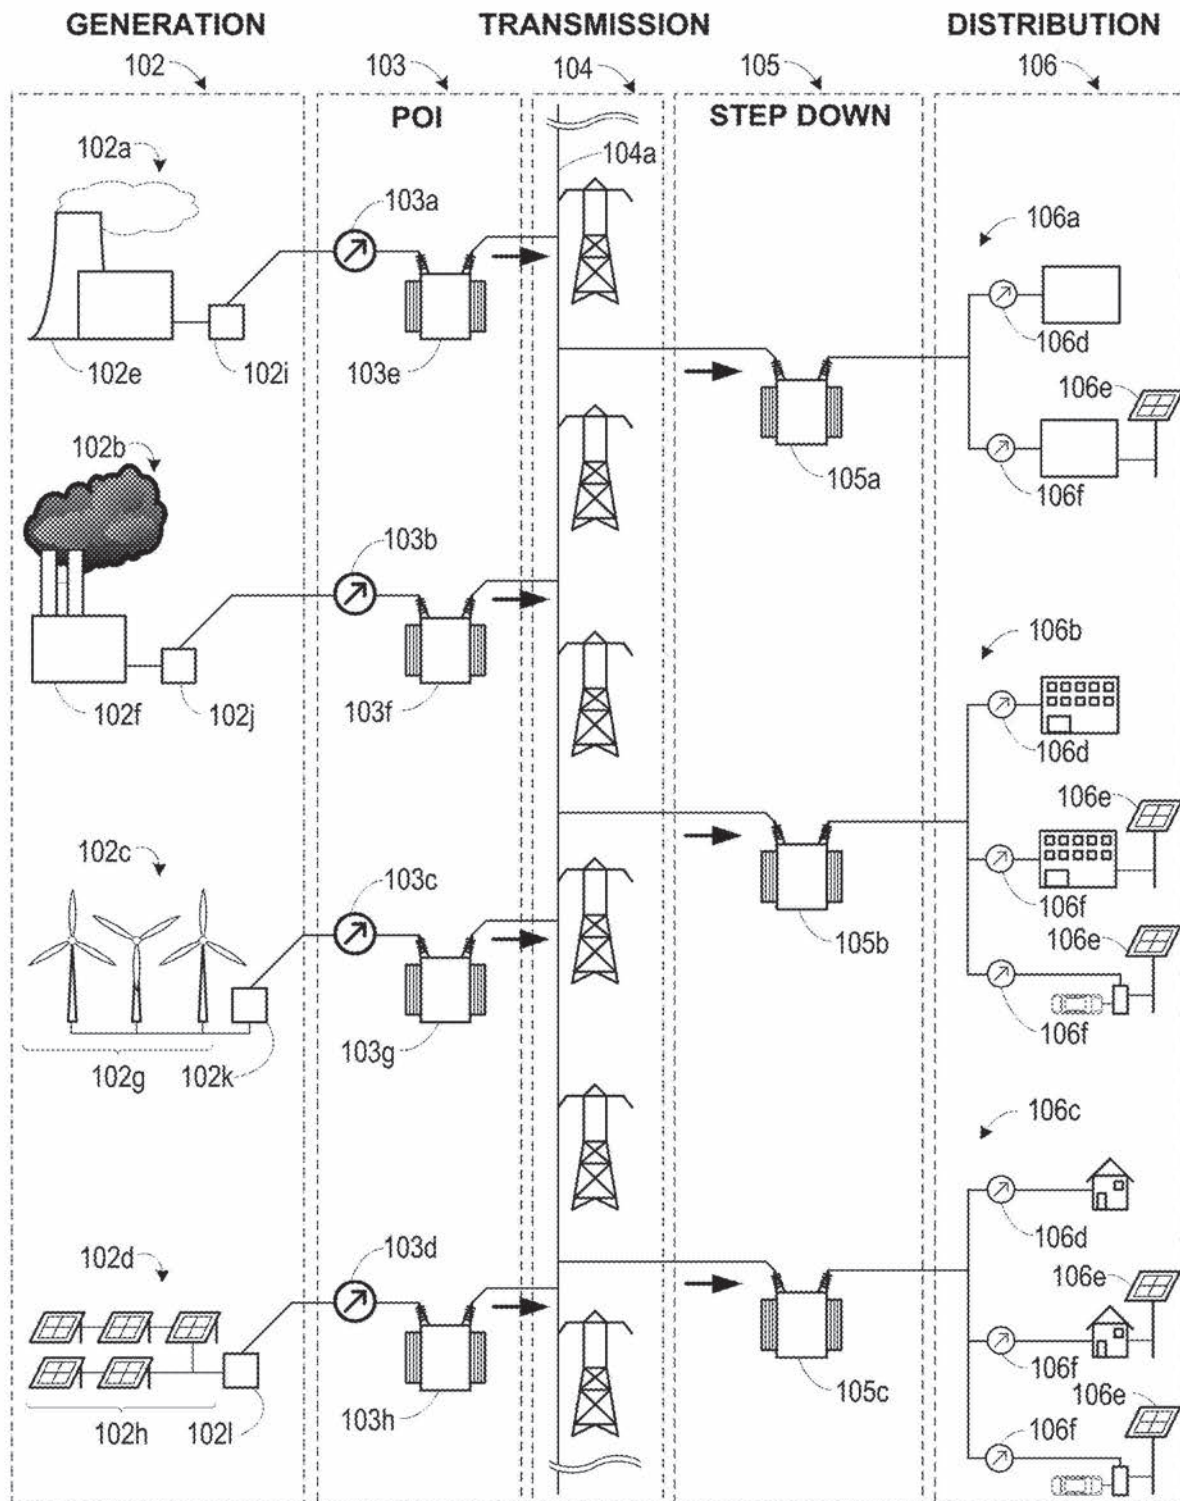
* cited by examiner

U.S. Patent

Mar. 31, 2020

Sheet 1 of 16

US 10,608,433 B1



PRIOR ART
FIGURE 1

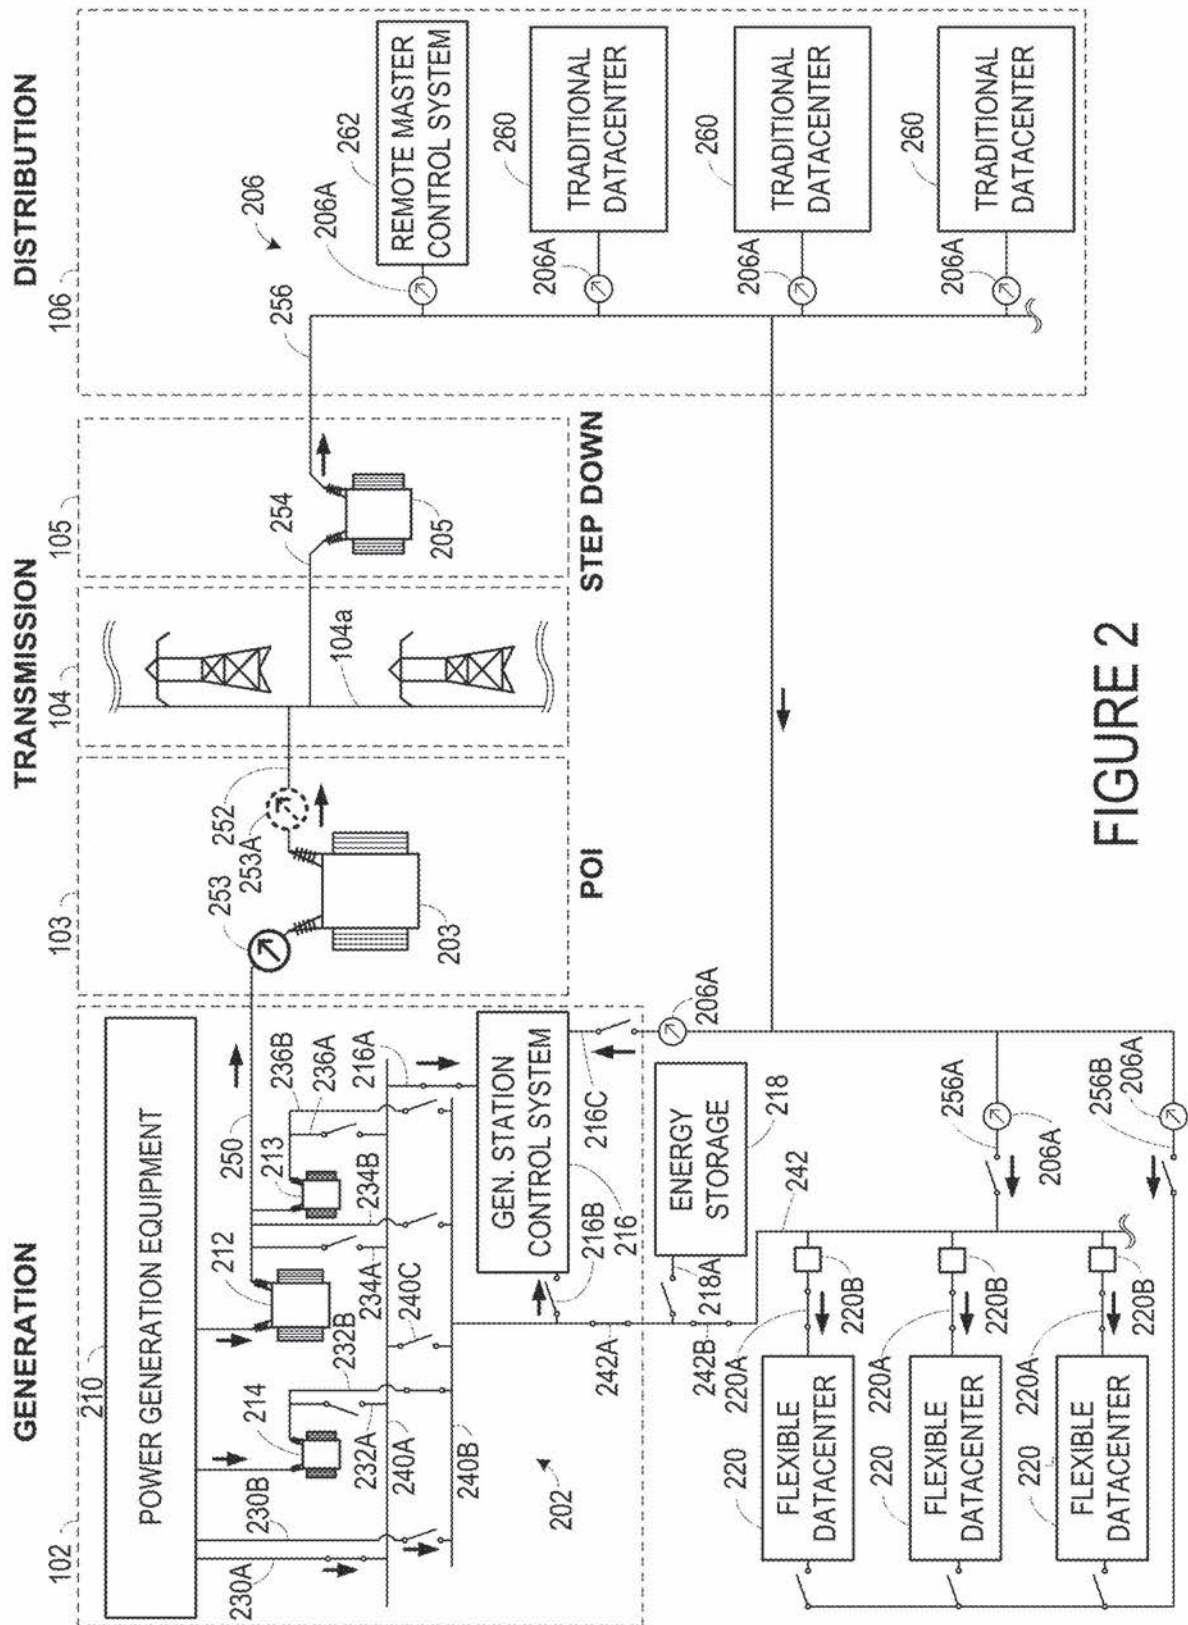


FIGURE 2

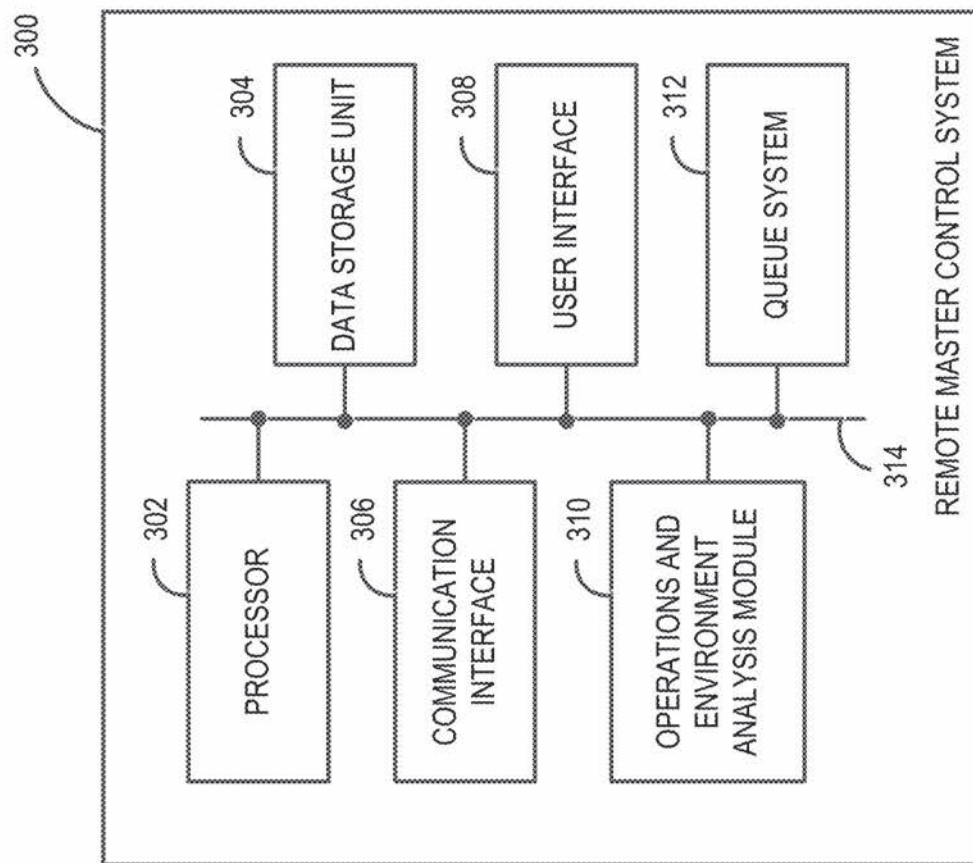


FIGURE 3

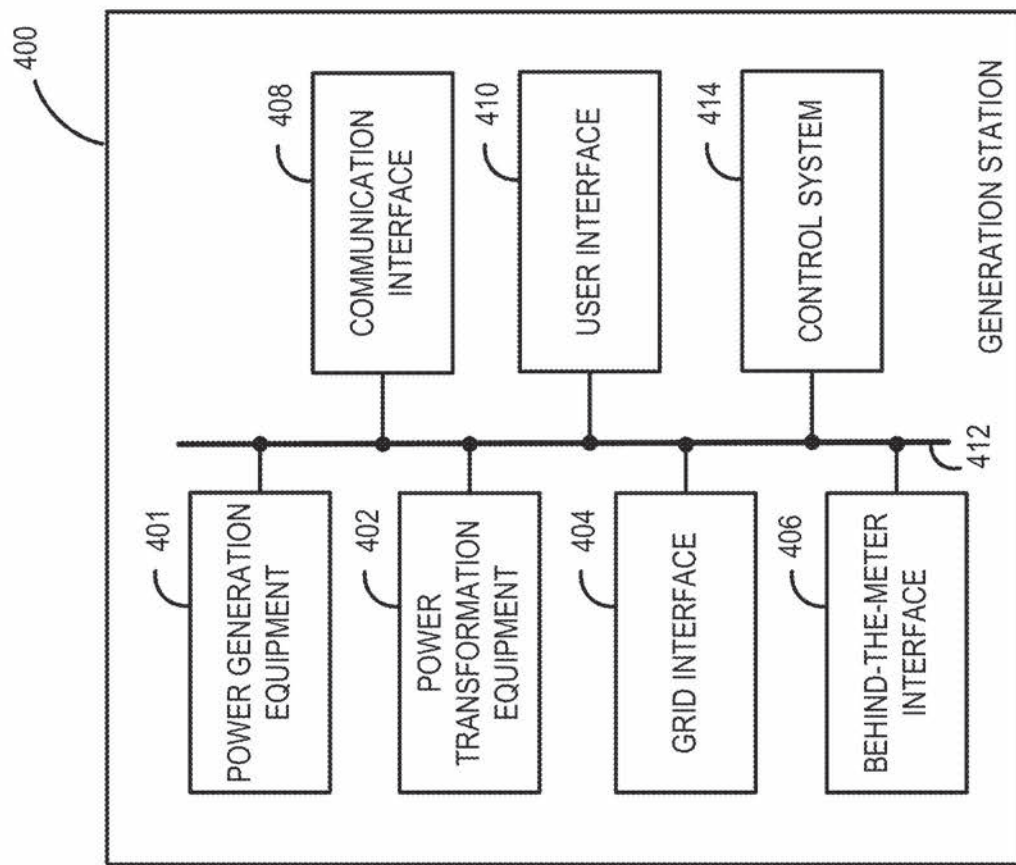


FIGURE 4

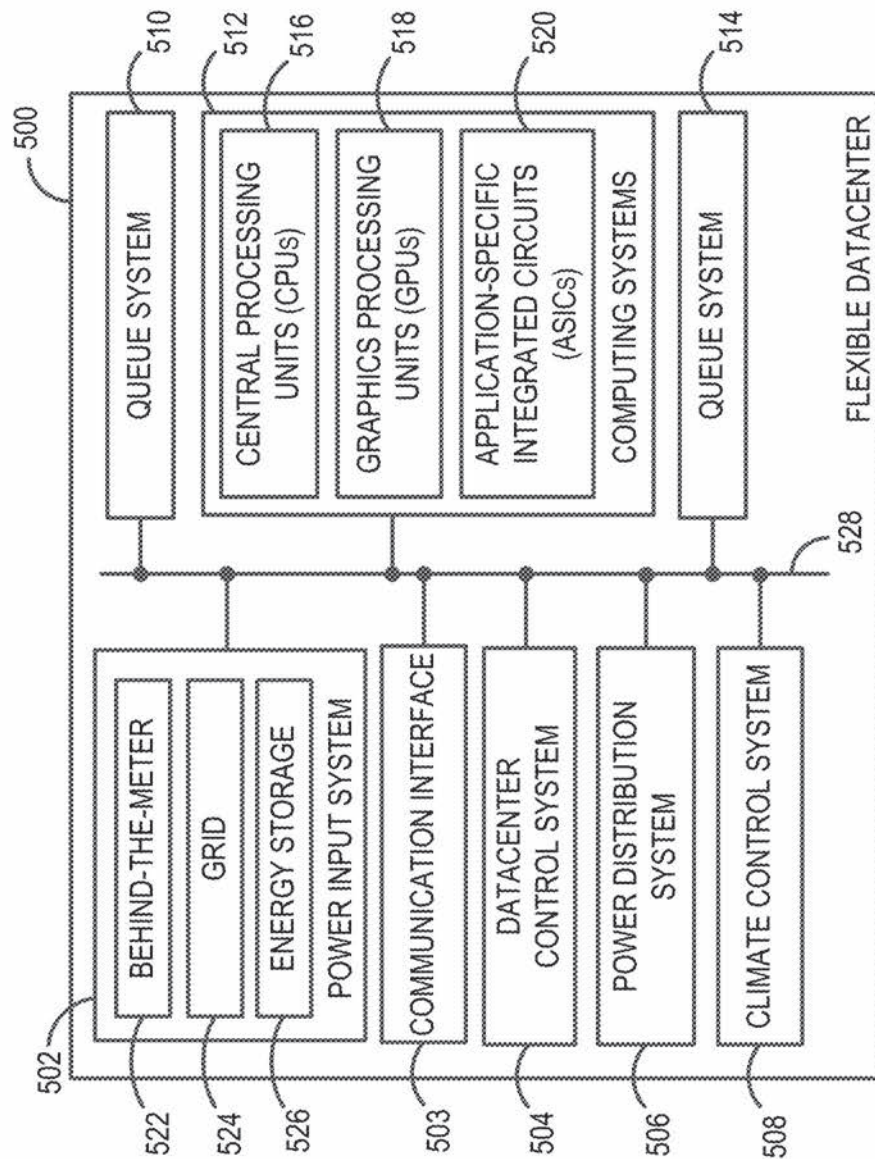


FIGURE 5

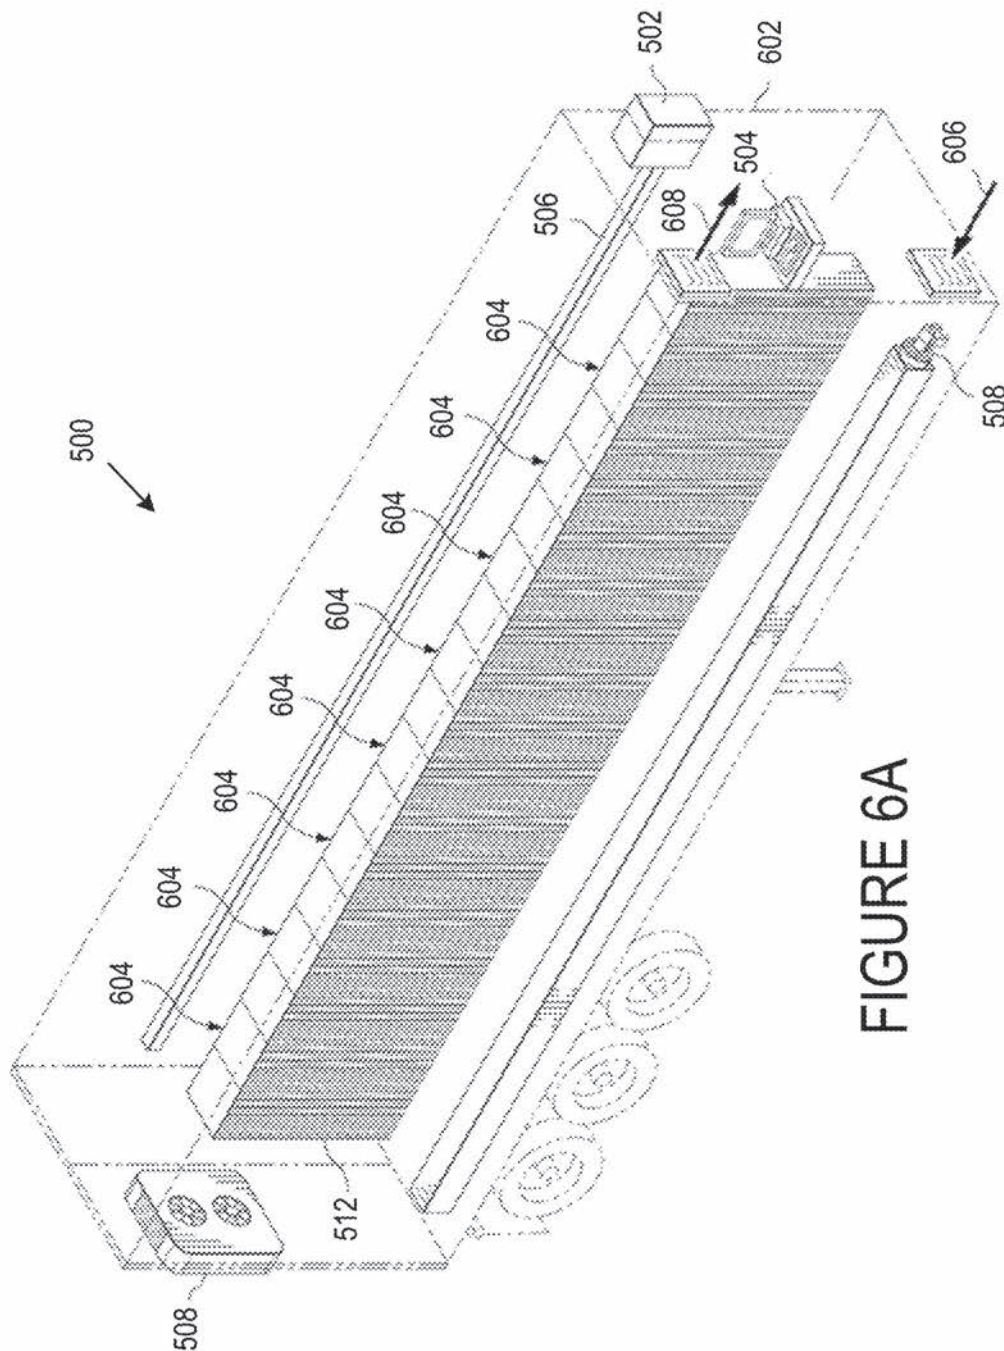


FIGURE 6A

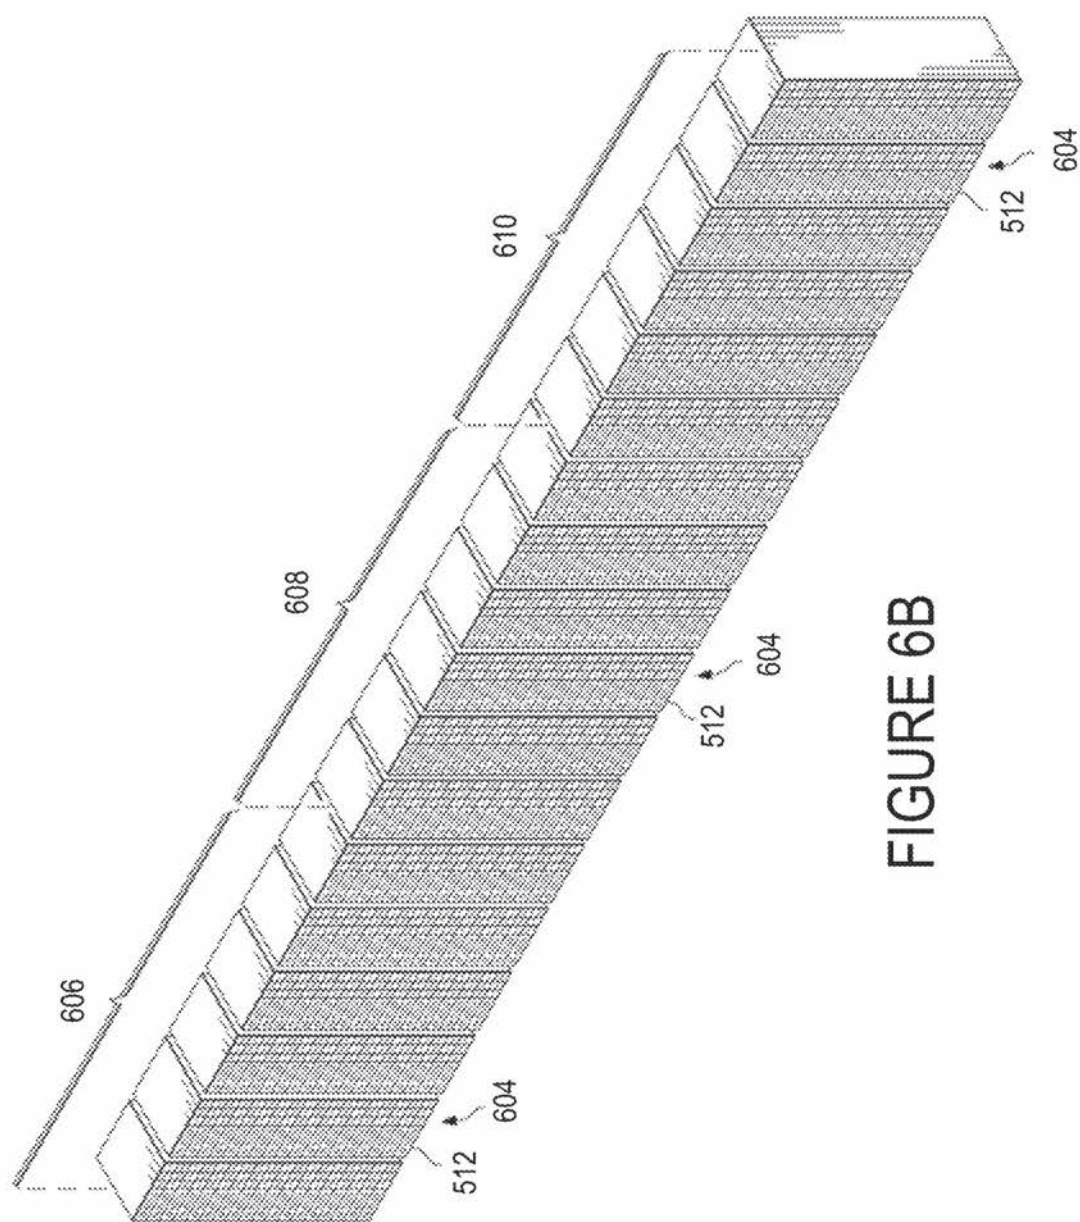


FIGURE 6B

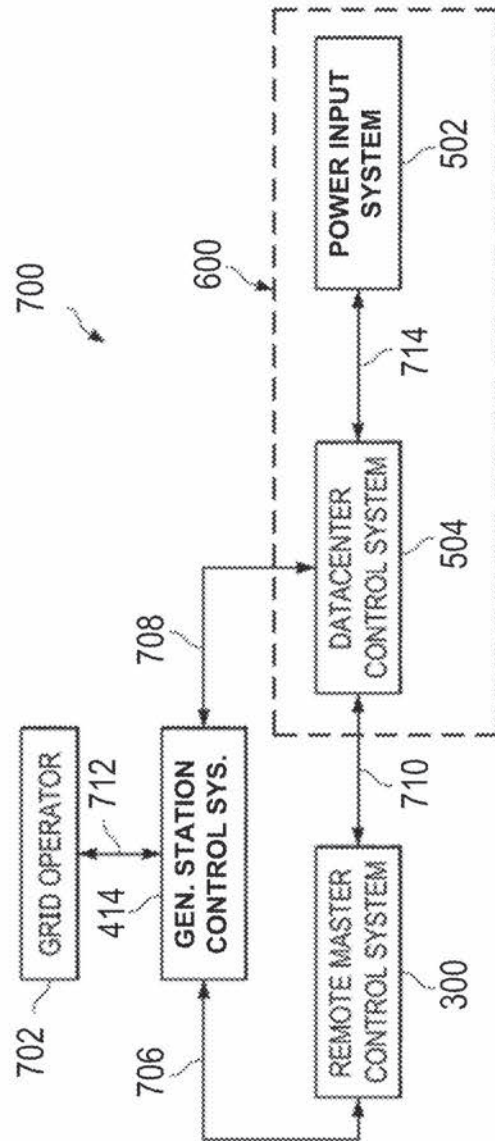
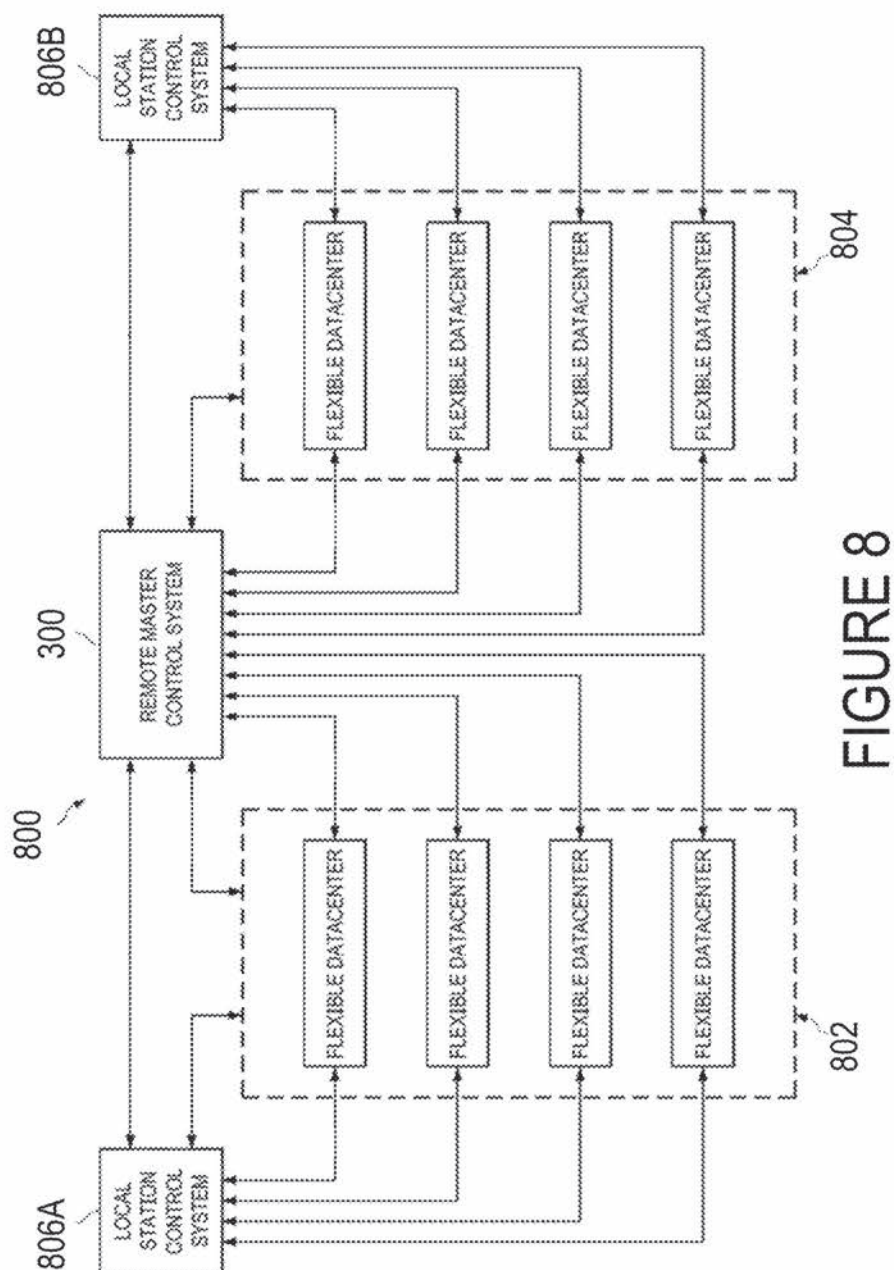


FIGURE 7



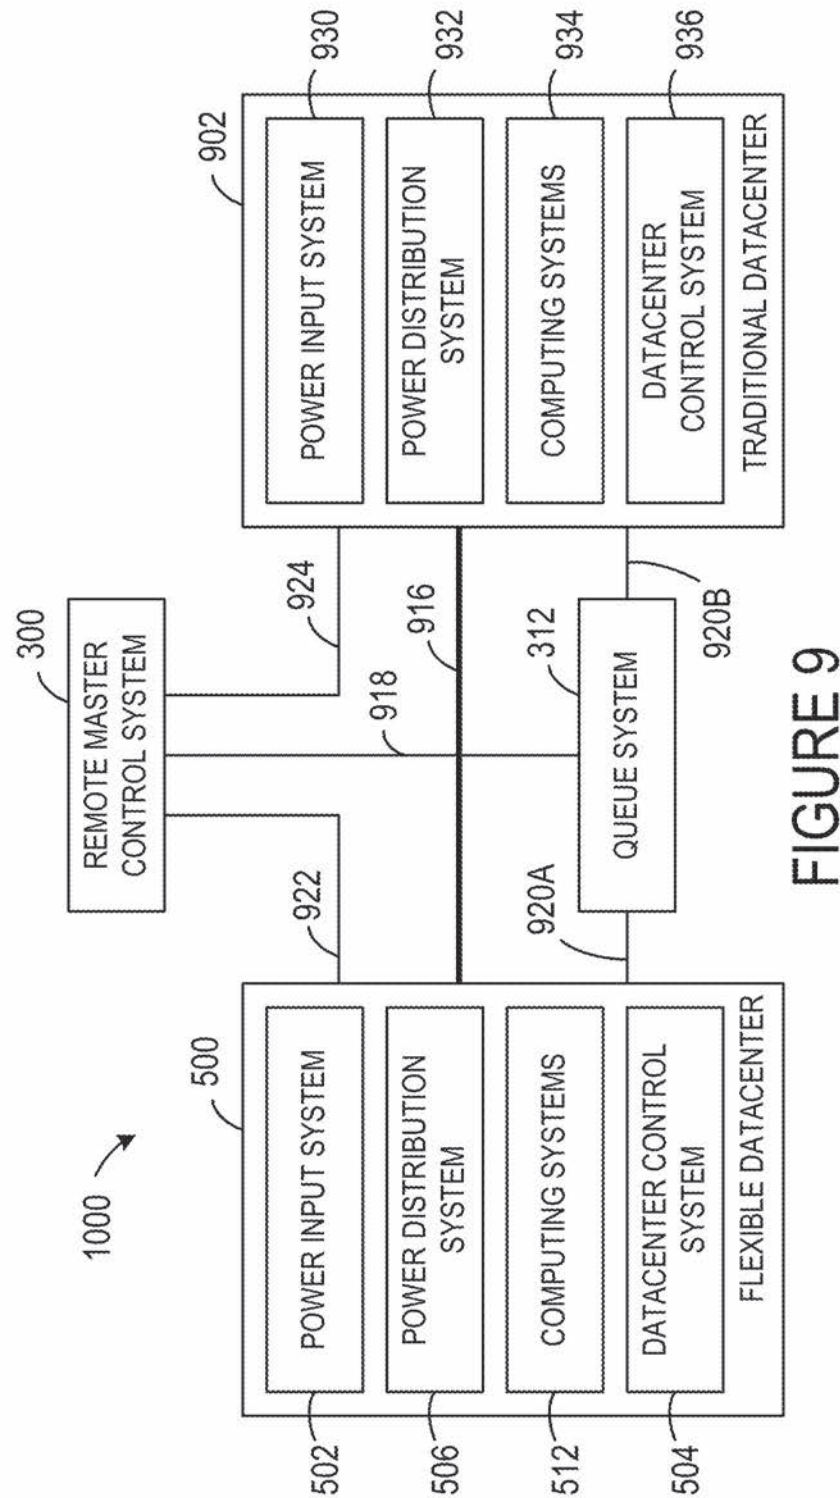


FIGURE 9

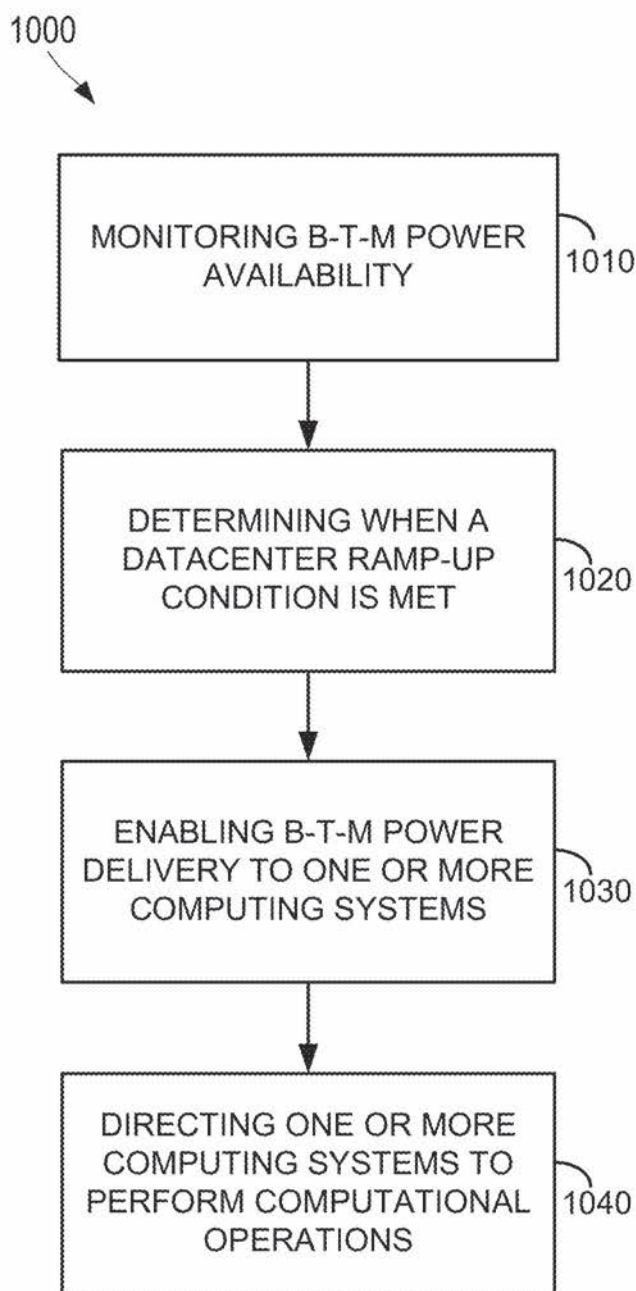


FIGURE 10A

U.S. Patent

Mar. 31, 2020

Sheet 12 of 16

US 10,608,433 B1

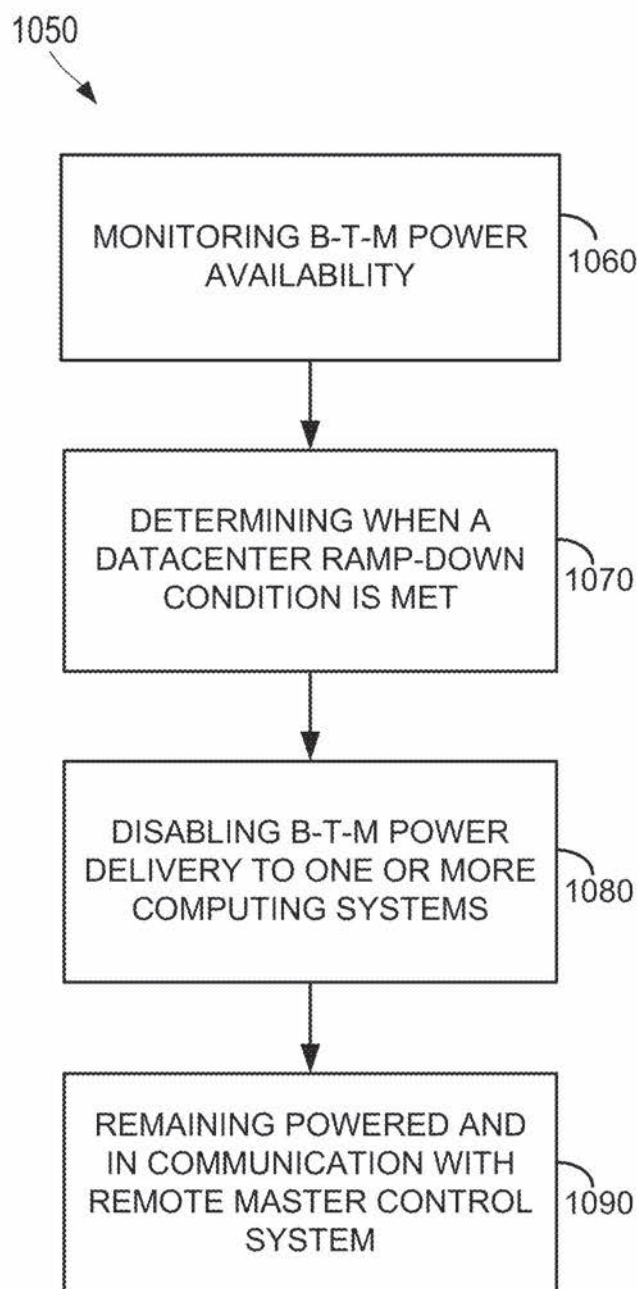


FIGURE 10B

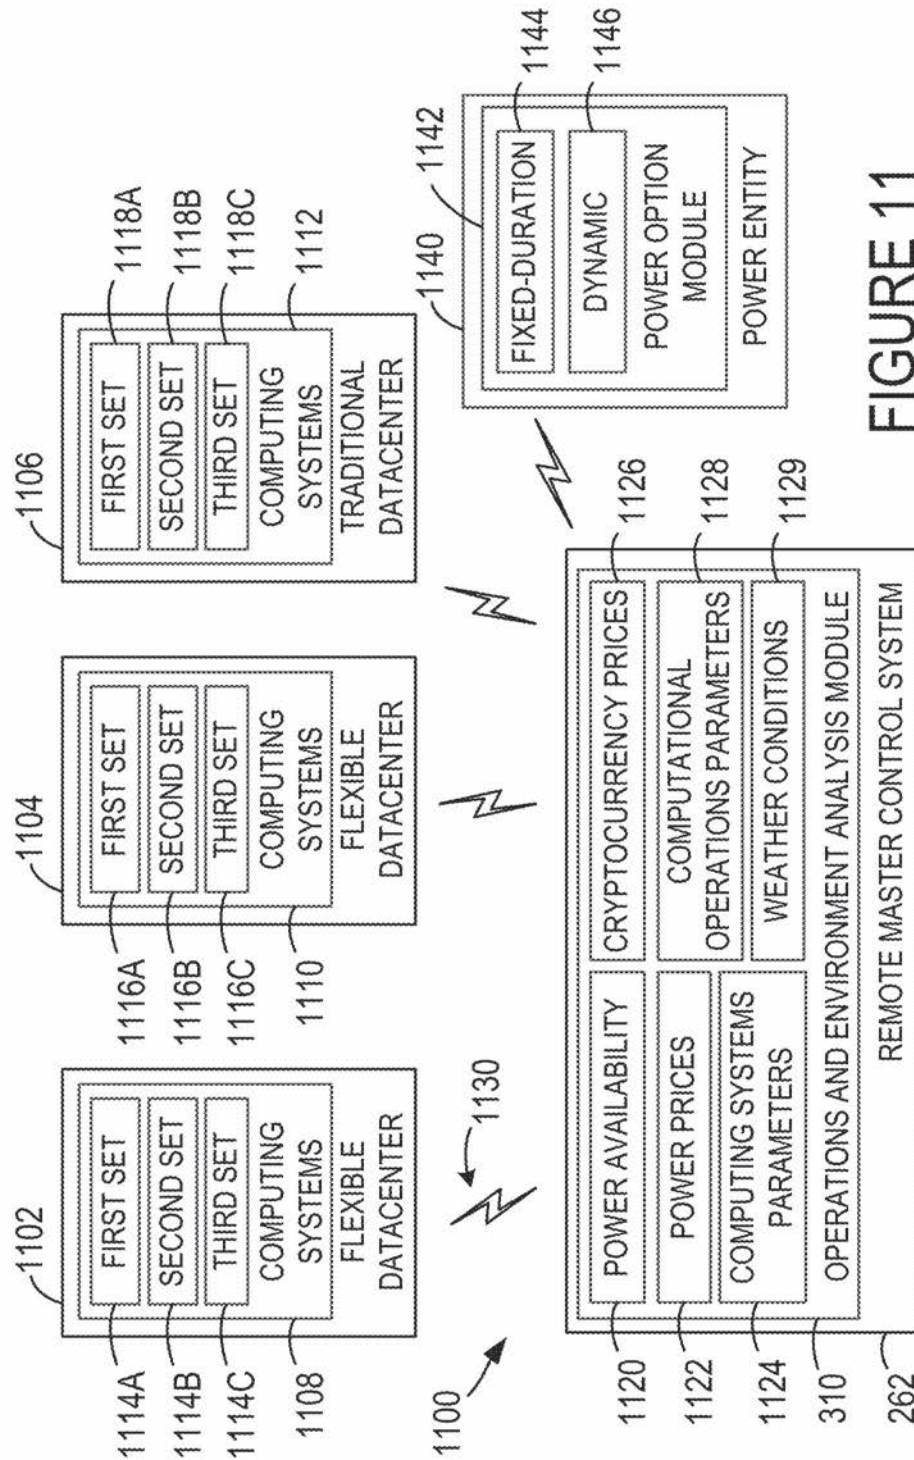


FIGURE 11

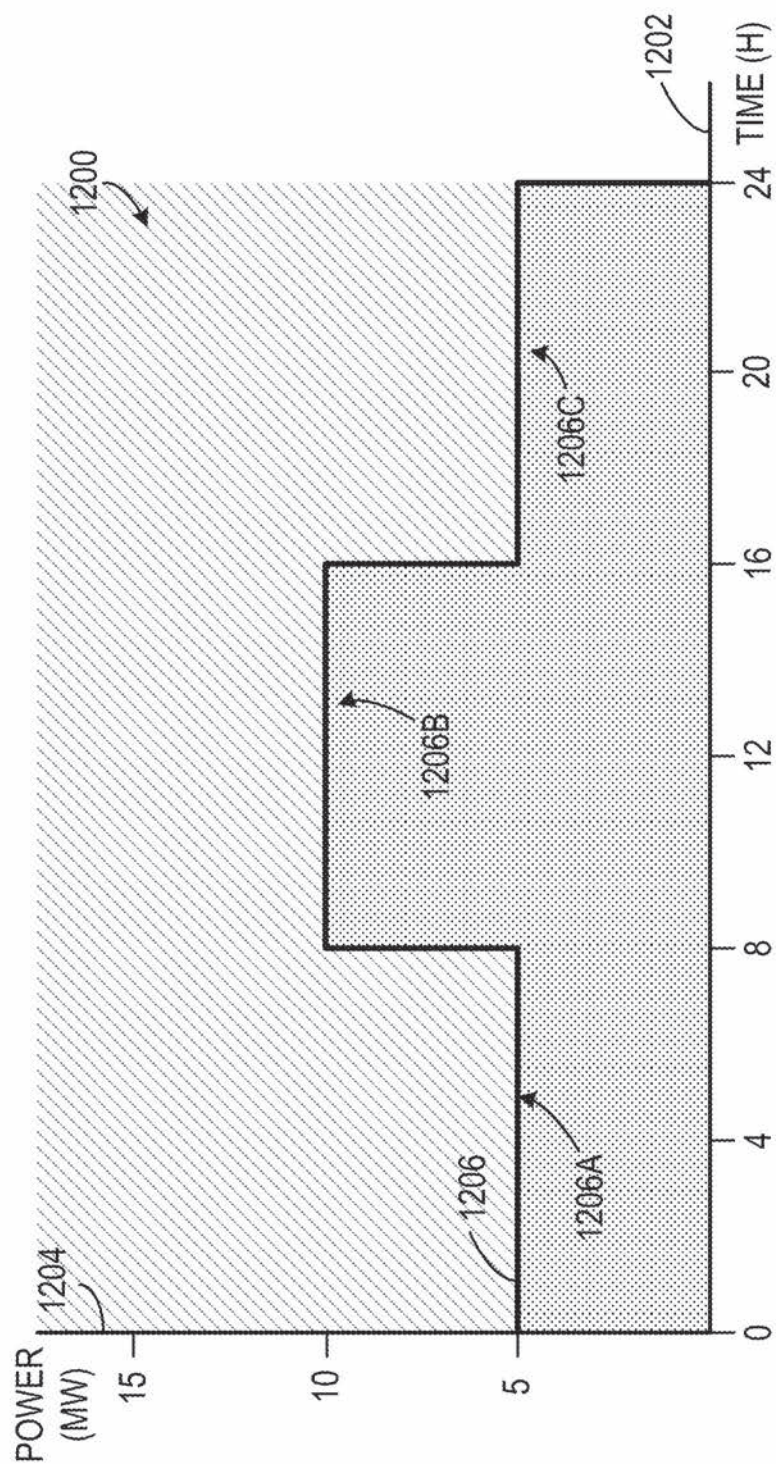


FIGURE 12

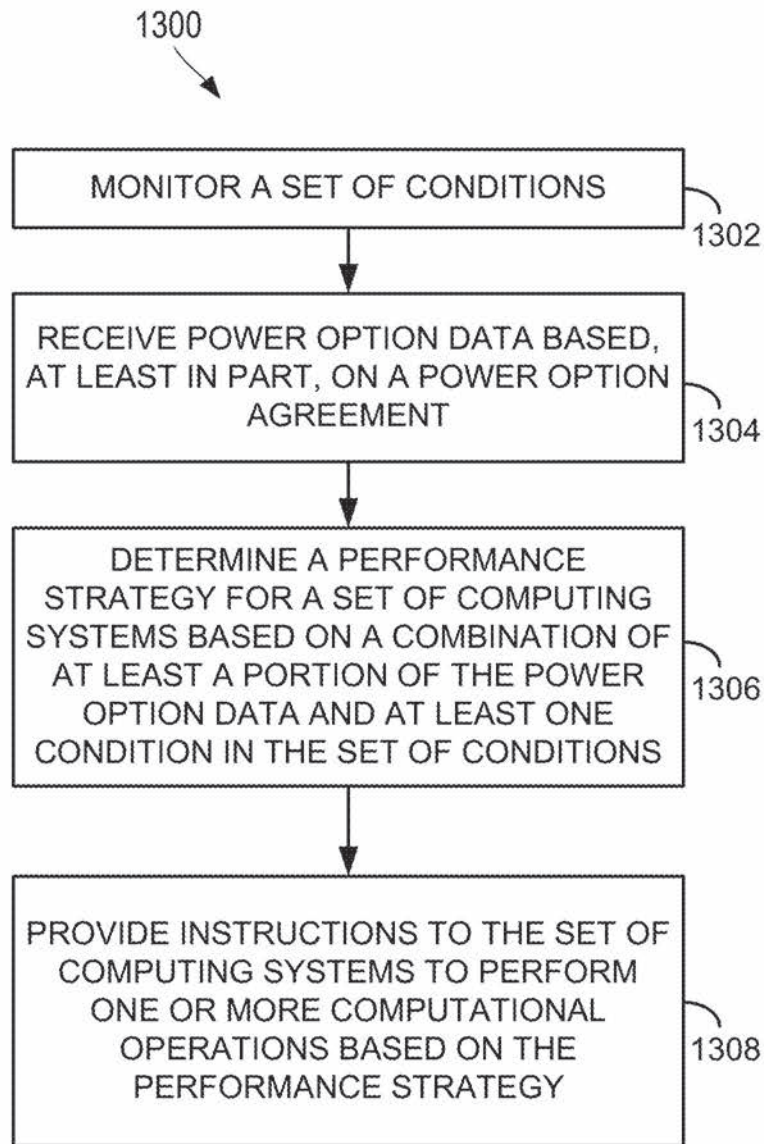


FIGURE 13

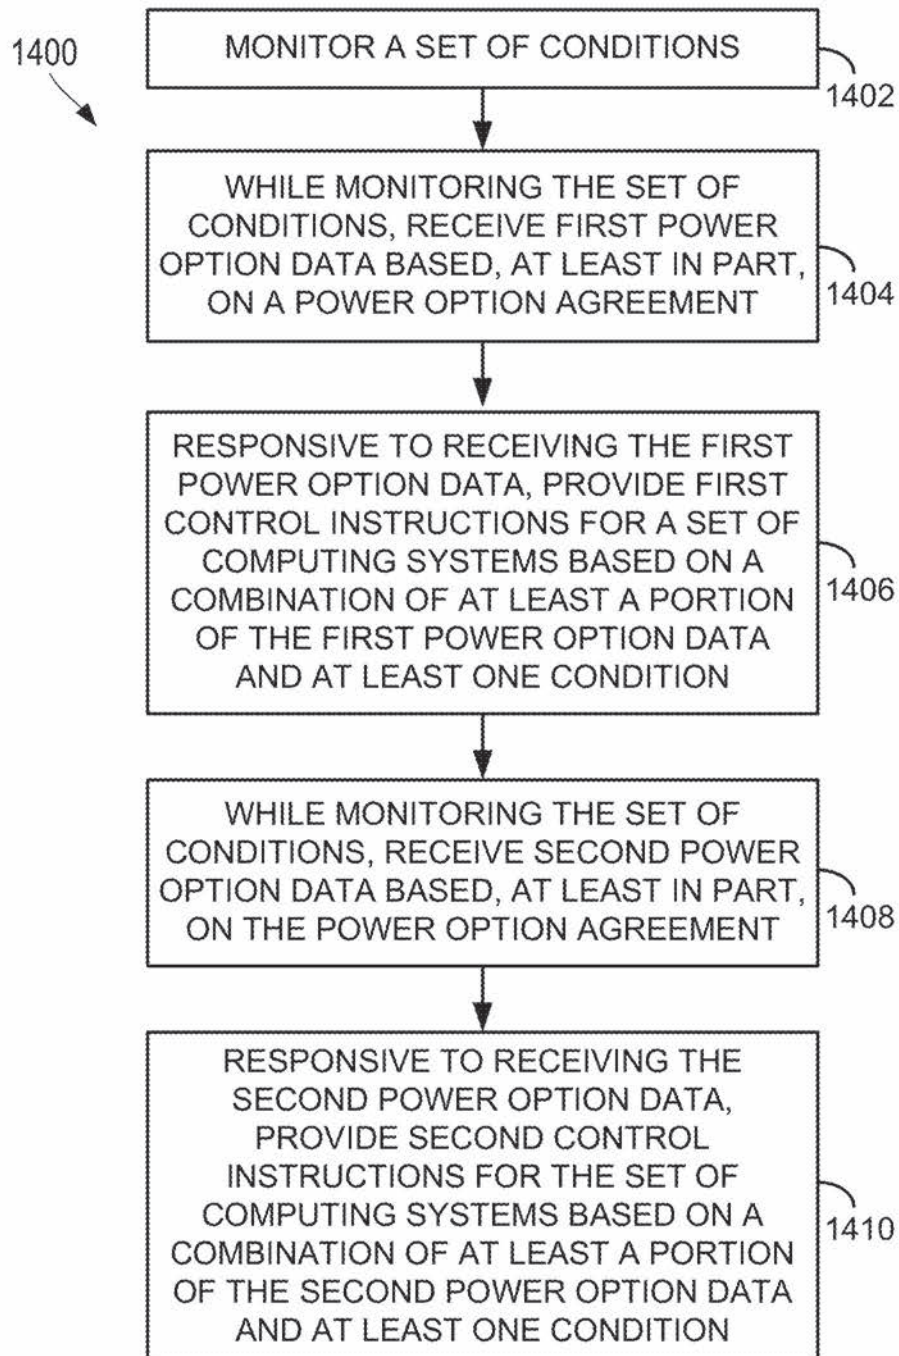


FIGURE 14

US 10,608,433 B1

1

METHODS AND SYSTEMS FOR ADJUSTING POWER CONSUMPTION BASED ON A FIXED-DURATION POWER OPTION AGREEMENT

CROSS-REFERENCE TO RELATED APPLICATIONS

The present application claims priority to U.S. Provisional Patent Application No. 62/927,119, filed Oct. 28, 2019, the entire contents of which are herein incorporated by reference.

FIELD

This specification relates to power consumption adjustments when using grid power and/or intermittent behind-the-meter power.

BACKGROUND

“Electrical grid” or “grid,” as used herein, refers to a Wide Area Synchronous Grid (also known as an Interconnection), and is a regional scale or greater electric power grid that operates at a synchronized frequency and is electrically tied together during normal system conditions. An electrical grid delivers electricity from generation stations to consumers. An electrical grid includes: (i) generation stations that produce electrical power at large scales for delivery through the grid, (ii) high voltage transmission lines that carry that power from the generation stations to demand centers, and (iii) distribution networks carry that power to individual customers.

FIG. 1 illustrates a typical electrical grid, such as a North American Interconnection or the synchronous grid of Continental Europe (formerly known as the UCTE grid). The electrical grid of FIG. 1 can be described with respect to the various segments that make up the grid.

A generation segment **102** includes one or more generation stations that produce utility-scale electricity (typically >50 MW), such as a nuclear plant **102a**, a coal plant **102b**, a wind power station (i.e., wind farm) **102c**, and/or a photovoltaic power station (i.e., a solar farm) **102d**. Generation stations are differentiated from building-mounted and other decentralized or local wind or solar power applications because they supply power at the utility level and scale (>50 MW), rather than to a local user or users. The primary purpose of generation stations is to produce power for distribution through the grid, and in exchange for payment for the supplied electricity. Each of the generation stations **102a-d** includes power generation equipment **102e-h**, respectively, typically capable of supply utility-scale power (>50 MW). For example, the power generation equipment **102g** at wind power station **102c** includes wind turbines, and the power generation equipment **102h** at photovoltaic power station **102d** includes photovoltaic panels.

Each of the generation stations **102a-d** may further include station electrical equipment **102i-1** respectively. Station electrical equipment **102i-1** are each illustrated in FIG. 1 as distinct elements for simplified illustrative purposes only and may, alternatively or additionally, be distributed throughout the power generation equipment, **102e-h**, respectively. For example, at wind power station **102c**, each wind turbine may include transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each wind turbine may be collected by distribution lines along strings of wind turbines and move through

2

collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the wind power station **102c**. Similarly, at photovoltaic power station **102d**, individual photovoltaic panels and/or arrays of photovoltaic panels may include inverters, transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each photovoltaic panel and/or array may be collected by distribution lines along the photovoltaic panels and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the photovoltaic power station **102d**.

Each generation station **102a-d** may produce AC or DC electrical current which is then typically stepped up to a higher AC voltage before leaving the respective generation station. For example, wind turbines may typically produce AC electrical energy at 600V to 700V, which may then be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102c**. As another example, photovoltaic arrays may produce DC voltage at 600V to 900V, which is then inverted to AC voltage and may be stepped up to 34.5 kV before leaving the generation station **102d**. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station **102d**.

Upon exiting the generation segment **102**, electrical power generated at generation stations **102a-d** passes through a respective Point of Interconnection (“POI”) **103** between a generation station (e.g., **102a-d**) and the rest of the grid. A respective POI **103** represents the point of connection between a generation station’s (e.g., **102a-d**) equipment and a transmission system (e.g., transmission segment **104**) associated with electrical grid. In some cases, at the POI **103**, generated power from generation stations **102a-d** may be stepped up at transformer systems **103e-h** to high voltage scales suitable for long-distance transmission along transmission lines **104a**. Typically, the generated electrical energy leaving the POI **103** will be at 115 kV AC or above, but in some cases it may be as low as, for example, 69 kV for shorter distance transmissions along transmission lines **104a**. Each of transformer systems **103e-h** may be a single transformer or may be multiple transformers operating in parallel or series and may be co-located or located in geographically distinct locations. Each of the transformer systems **103e-h** may include substations and other links between the generation stations **102a-d** and the transmission lines **104a**.

A key aspect of the POI **103** is that this is where generation-side metering occurs. One or more utility-scale generation-side meters **103a-d** (e.g., settlement meters) are located at settlement metering points at the respective POI **103** for each generation station **102a-d**. The utility-scale generation-side meters **103a-d** measure power supplied from generation stations **102a-d** into the transmission segment **104** for eventual distribution throughout the grid.

For electricity consumption, the price consumers pay for power distributed through electric power grids is typically composed of, among other costs, Generation, Administration, and Transmission & Distribution (“T&D”) costs. T&D costs represent a significant portion of the overall price paid by consumers for electricity. These costs include capital costs (land, equipment, substations, wire, etc.), costs associated with electrical transmission losses, and operation and maintenance costs.

US 10,608,433 B1

3

For utility-scale electricity supply, operators of generation stations (e.g., 102a-d) are paid a variable market price for the amount of power the operator generates and provides to the grid, which is typically determined via a power purchase agreement (PPA) between the generation station operator and a grid operator. The amount of power the generation station operator generates and provides to the grid is measured by utility-scale generation-side meters (e.g., 103a-d) at settlement metering points. As illustrated in FIG. 1, the utility-scale generation-side meters 103a-d are shown on a low side of the transformer systems 103e-h, but they may alternatively be located within the transformer systems 103e-h or on the high side of the transformer systems 103e-h. A key aspect of a utility-scale generation-side meter is that it is able to meter the power supplied from a specific generation station into the grid. As a result, the grid operator can use that information to calculate and process payments for power supplied from the generation station to the grid. That price paid for the power supplied from the generation station is then subject to T&D costs, as well as other costs, in order to determine the price paid by consumers.

After passing through the utility-scale generation-side meters in the POI 103, the power originally generated at the generation stations 102a-d is transmitted onto and along the transmission lines 104a in the transmission segment 104. Typically, the electrical energy is transmitted as AC at 115 kV+ or above, though it may be as low as 69 kV for short transmission distances. In some cases, the transmission segment 104 may include further power conversions to aid in efficiency or stability. For example, transmission segment 104 may include high-voltage DC ("HVDC") portions (along with conversion equipment) to aid in frequency synchronization across portions of the transmission segment 104. As another example, transmission segment 104 may include transformers to step AC voltage up and then back down to aid in long distance transmission (e.g., 230 kV, 500 kV, 765 kV, etc.).

Power generated at the generation stations 104a-d is ultimately destined for use by consumers connected to the grid. Once the energy has been transmitted along the transmission segment 104, the voltage will be stepped down by transformer systems 105a-c in the step down segment 105 so that it can move into the distribution segment 106.

In the distribution segment 106, distribution networks 106a-c take power that has been stepped down from the transmission lines 104a and distribute it to local customers, such as local sub-grids (illustrated at 106a), industrial customers, including large EV charging networks (illustrated at 106b), and/or residential and retail customers, including individual EV charging stations (illustrated at 106c). Customer meters 106d, 106f measure the power used by each of the grid-connected customers in distribution networks 106a-c. Customer meters 106d are typically load meters that are unidirectional and measure power use. Some of the local customers in the distribution networks 106a-d may have local wind or solar power systems 106e owned by the customer. As discussed above, these local customer power systems 106e are decentralized and supply power directly to the customer(s). Customers with decentralized wind or solar power systems 106e may have customer meters 106f that are bidirectional or net-metering meters that can track when the local customer power systems 106e produce power in excess of the customer's use, thereby allowing the utility to provide a credit to the customer's monthly electricity bill. Customer meters 106d, 106f differ from utility-scale generation-side meters (e.g., settlement meters) in at least the following characteristics: design (electro-mechanical or electronic vs

4

current transformer), scale (typically less than 1600 amps vs. typically greater than 50 MW; typically less than 600V vs. typically greater than 14 kV), primary function (use vs. supply metering), economic purpose (credit against use vs. payment for power), and location (in a distribution network at point of use vs. at a settlement metering point at a Point of Interconnection between a generation station and a transmission line).

To maintain stability of the grid, the grid operator strives to maintain a balance between the amount of power entering the grid from generation stations (e.g., 102a-d) and the amount of grid power used by loads (e.g., customers in the distribution segment 106). In order to maintain grid stability and manage congestion, grid operators may take steps to reduce the supply of power arriving from generation stations (e.g., 102a-d) when necessary (e.g., curtailment). Particularly, grid operators may decrease the market price paid for generated power to dis-incentivize generation stations (e.g., 102a-d) from generating and supplying power to the grid. In some cases, the market price may even go negative such that generation station operators must pay for power they allow into the grid. In addition, some situations may arise where grid operators explicitly direct a generation station (e.g., 102a-d) to reduce or stop the amount of power the station is supplying to the grid.

Power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewable energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

According to the National Renewable Energy Laboratory's Technical Report TP-6A20-60983 (March 2014):

[C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission congestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

"Curtailment" today does not necessarily mean what it did in the early 2000s. Two separate changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them. Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being

US 10,608,433 B1

5

replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator’s capital costs are included in a utility’s rate base.

Curtailment may result in available energy being wasted because solar and wind operators have zero variable cost (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

SUMMARY

In an example, a system includes a set of computing systems. The set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and, while monitoring the set of conditions, receive first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The control system is further configured to provide first control instructions for the set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions responsive to receiving the first power option data. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The control system is also configured to, while monitoring the set of conditions, receive second power option data based, at least in part, on the power option

6

agreement. The second power option data specify a second minimum power threshold associated with a second time interval. Responsive to receiving the second power option data, the control system is configured to provide second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and wherein the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In another example, a method involves monitoring, at a computing system, a set of conditions, and while monitoring the set of conditions, receiving first power option data based, at least in part, on a power option agreement. The first power option data specify a first minimum power threshold associated with a first time interval. The method further involves, responsive to receiving the first power option data, providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition of the set of conditions. The first control instructions comprises a first power consumption target for the set of computing systems for the first time interval, and the first power consumption target is equal to or greater than the first minimum power threshold associated with the first time interval. The method further involves, while monitoring the set of conditions, receiving second power option data based, at least in part, on the power option agreement. The second power option data specify a second minimum power threshold associated with a second time interval. The method also involves, responsive to receiving the second power option data, providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power data and at least one condition of the set of conditions. The second control instructions comprises a second power consumption target for the set of computing systems for the second time interval, and the second power consumption target is equal to or greater than the second minimum power threshold associated with the second time interval.

In yet another example, a system is provided. The system includes a set of computing systems, where the set of computing systems is configured to perform computational operations using power from a power grid. The system also includes a control system configured to monitor a set of conditions and receive power option data based, at least in part, on a power option agreement. The power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, where each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals. The control system is further configured to, responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, where each power consumption target is equal to or greater than the minimum power threshold associated with each time interval. The control system is also configured to provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.

In a further example, non-transitory computer-readable medium is described that is configured to store instructions,

that when executed by a computing system, causes the computing system to perform operations consistent with the method steps described above.

Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE FIGURES

FIG. 1 shows a typical electrical grid.

FIG. 2 shows a behind-the-meter arrangement with optional grid power, including one or more flexible datacenters, according to one or more example embodiments.

FIG. 3 shows a block diagram of a remote master control system, according to one or more example embodiments.

FIG. 4 a block diagram of a generation station, according to one or more example embodiments.

FIG. 5 shows a block diagram of a flexible datacenter, according to one or more example embodiments.

FIG. 6A shows a structural arrangement of a flexible datacenter, according to one or more example embodiments.

FIG. 6B shows a set of computing systems arranged in a straight configuration, according to one or more example embodiments.

FIG. 7 shows a control distribution system for a flexible datacenter, according to one or more example embodiments.

FIG. 8 shows a control distribution system for a fleet of flexible datacenters, according to one or more example embodiments.

FIG. 9 shows a queue distribution system for a traditional datacenter and a flexible datacenter, according to one or more example embodiments.

FIG. 10A shows a method of dynamic power consumption at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 10B shows a method of dynamic power delivery at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

FIG. 11 shows a block diagram of a system for implementing power consumption adjustments based on a power option agreement, according to one or more embodiments.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments.

FIG. 13 shows a method for implementing power consumption adjustments based on a fixed-duration power option agreement, according to one or more embodiments.

FIG. 14 shows a method for implementing power consumption adjustments based on a dynamic power option agreement, according to one or more embodiments.

DETAILED DESCRIPTION

Disclosed examples will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all of the disclosed examples are shown. Different examples may be described and should not be construed as limited to the examples set forth herein.

As discussed above, the market price paid to generation stations for supplying power to the grid often fluctuates due to various factors, including the need to maintain grid stability and based on current demand and usage by connected loads in distribution networks. Due to these factors, situations can arise where generation stations are offered substantially lower prices to deter an over-supply of power to the grid. Although these situations typically exist temporarily, generation stations are sometimes forced to either sell power to the grid at the much lower prices or adjust

operations to decrease the amount of power generated. Furthermore, some situations may even require generation stations to incur costs in order to offload power to the grid or to shut down generation temporarily.

The volatility in the market price offered for power supplied to the grid can be especially problematic for some types of generation stations. In particular, wind farms and some other types of renewable resource power producers may lack the ability to quickly adjust operations in response to changes in the market price offered for supplying power to the grid. As a result, power generation and management at some generation stations can be inefficient, which can frequently result in power being sold to the grid at low or negative prices. In some situations, a generation station may even opt to halt power generation temporarily to avoid such unfavorable pricing. As such, the time required to halt and to restart the power generation at a generation station can reduce the generation station's ability to take advantage of rising market prices for power supplied to the grid.

Example embodiments provided herein aim to assist generation stations in managing power generation operations and avoid unfavorable power pricing situations like those described above. In particular, example embodiments may involve providing a load that is positioned behind-the-meter ("BTM") and enabling the load to utilize power received behind-the-meter at a generation station in a timely manner. As a general rule of thumb, BTM power is not subject to traditional T&D costs.

For purposes herein, a generation station is considered to be configured for the primary purpose of generating utility-scale power for supply to the electrical grid (e.g., a Wide Area Synchronous Grid or a North American Interconnect).

In one embodiment, equipment located behind-the-meter ("BTM equipment") is equipment that is electrically connected to a generation station's power generation equipment behind (i.e., prior to) the generation station's POI with an electrical grid.

In one embodiment, behind-the-meter power ("BTM power") is electrical power produced by a generation station's power generation equipment and utilized behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM equipment receives power from the generation station, but the power received by the BTM equipment from the generation station has not passed through the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM power is utilized before being metered at the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station—for example, a substation between the generation station and the generation station's POI.

US 10,608,433 B1

9

In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that supplies power to a grid, and the BTM equipment receives power from the generation station that is not subject to T&D charges, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that supplies power to a grid, and the BTM power is not subject to T&D charges before being used by electrical equipment, but power received from the grid that is supplied by the generation station is subject to T&D charges.

In another embodiment, equipment may be considered behind-the-meter if the BTM equipment receives power generated from the generation station and that received power is not routed through the electrical grid before being delivered to the BTM equipment.

In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station, and BTM equipment receives that generated power, and that generated power received by the BTM equipment is not routed through the electrical grid before being delivered to the BTM equipment.

For purposes herein, BTM equipment may also be referred to as a behind-the-meter load ("BTM load") when the BTM equipment is actively consuming BTM power.

Beneficially, where BTM power is not subject to traditional T&D costs, a wind farm or other type of generation station can be connected to BTM loads which can allow the generation station to selectively avoid the adverse or less-than optimal cost structure occasionally associated with supplying power to the grid by shunting generated power to the BTM load.

An arrangement that positions and connects a BTM load to a generation station can offer several advantages. In such arrangements, the generation station may selectively choose whether to supply power to the grid or to the BTM load, or both. The operator of a BTM load may pay to utilize BTM power at a cost less than that charged through a consumer meter (e.g., 106d, 1060 located at a distribution network (e.g., 106a-c) receiving power from the grid. The operator of a BTM load may additionally or alternatively charge less than the market rate to consume excess power generated at the generation station during curtailment. As a result, the generation station may direct generated power based on the "best" price that the generation station can receive during a given time frame, and/or the lowest cost the generation station may incur from negative market pricing during curtailment. The "best" price may be the highest price that the generation station may receive for its generated power during a given duration, but can also differ within embodiments and may depend on various factors, such as a prior PPA.

In one example, by having a behind-the-meter option available, a generation station may transition from supplying all generated power to the grid to supplying some or all generated power to one or more BTM loads when the market price paid for power by grid operators drops below a predefined threshold (e.g., the price that the operator of the BTM load is willing to pay the generation station for power). Thus, by having an alternative option for power consumption (i.e., one or more BTM loads), the generation station can selectively utilize the different options to maximize the price received for generated power. In addition, the generation station may also utilize a BTM load to avoid or reduce

10

the economic impact in situations when supplying power to the grid would result in the generation station incurring a net cost.

Providing BTM power to a load can also benefit the BTM load operator. A BTM load may be able to receive and utilize BTM power received from the generation station at a cost that is lower than the cost for power from the grid (e.g., at a customer meter 106d, 1060. This is primarily due to the avoidance (or significant reduction) in T&D costs and the market effects of curtailment. As indicated above, the generation station may be willing to divert generated power to the BTM load rather than supplying the grid due to changing market conditions, or during maintenance periods, or for other non-market conditions. Thus, some situations may arise where the generation station offers power to the BTM load at a price that is substantially lower than the price available on the grid. Furthermore, in some situations, the BTM load may even be able to obtain and utilize BTM power from a generation station at no cost or even at negative pricing since the generation station may rather supply the BTM load with generated power during a given time range instead of paying a higher price for the grid to take the power or modifying operations to decrease power output.

Another example of cost-effective use of BTM power is when the generation station 202 is selling power to the grid at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price the generation station 202 would have to pay to the grid power to offload generation's station 202 generated power. Advantageously, one or more flexible datacenters 220 may take the generated power behind-the-meter, thereby allowing the generation station 202 to produce and obtain the production tax credit, while selling less power to the grid at the negative price.

Another example of cost-effective behind-the-meter power is when the generation station 202 is selling power to the grid at a negative price because the grid is oversupplied and/or the generation station 202 is instructed to stand down and stop producing altogether. A grid operator may select and direct certain generation stations to go offline and stop supplying power to the grid. Advantageously, one or more flexible datacenters may be used to take power behind-the-meter, thereby allowing the generation station 202 to stop supplying power to the grid, but still stay online and make productive use of the power generated.

Another example of beneficial behind-the-meter power use is when the generation station 202 is producing power that is, with reference to the grid, unstable, out of phase, or at the wrong frequency, or the grid is already unstable, out of phase, or at the wrong frequency. A grid operator may select certain generation stations to go either offline and stop producing power, or to take corrective action with respect to the grid power stability, phase, or frequency. Advantageously, one or more flexible datacenters 220 may be used to selectively consume power behind-the-meter, thereby allowing the generation station 202 to stop providing power to the grid and/or provide corrective feedback to the grid.

Another example of beneficial behind-the-meter power use is that cost-effective behind-the-meter power availability may occur when the generation station 202 is starting up or testing. Individual equipment in the power generation equipment 210 may be routinely offline for installation, maintenance, and/or service and the individual units must be tested prior to coming online as part of overall power generation equipment 210. During such testing or maintenance time, one or more flexible datacenters may be intermittently

US 10,608,433 B1

11

powered by the one or more units of the power generation equipment **210** that are offline from the overall power generation equipment **210**.

Another example of beneficial behind-the-meter power use is that datacenter control systems at the flexible datacenters **220** may quickly ramp up and ramp down power consumption by computing systems in the flexible datacenters **220** based on power availability from the generation station **202**. For instance, if the grid requires additional power and signals the demand via a higher local price for power, the generation station **202** can supply the grid with power nearly instantly by having active flexible datacenters **220** quickly ramp down and turn off computing systems (or switch to a stored energy source), thereby reducing an active BTM load.

Another example of beneficial behind-the-meter power use is in new photovoltaic generation stations **202**. For example, it is common to design and build new photovoltaic generation stations with a surplus of power capacity to account for degradation in efficiency of the photovoltaic panels over the life of the generation stations. Excess power availability at the generation station can occur when there is excess local power generation and/or low grid demand. In high incident sunlight situations, a photovoltaic generation station **202** may generate more power than the intended capacity of generation station **202**. In such situations, a photovoltaic generation station **202** may have to take steps to protect its equipment from damage, which may include taking one or more photovoltaic panels offline or shunting their voltage to dummy loads or the ground. Advantageously, one or more flexible datacenters (e.g., the flexible datacenters **220**) may take power behind-the-meter at the Generation Station **202**, thereby allowing the generation station **202** to operate the power generation equipment **210** within operating ranges while the flexible datacenters **220** receive BTM power without transmission or distribution costs.

Thus, for at least the reasons described herein, arrangements that involves providing a BTM load as an alternative option for a generation station to direct its generated power to can serve as a mutually beneficial relationship in which both the generation station and the BTM load can economically benefit. The above-noted examples of beneficial use of BTM power are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as benefits to unutilized BTM power capacity, BTM power pricing, or BTM power consumption.

Within example embodiments described herein, various types of utility-scale power producers may operate as generation stations **202** that are capable of supplying power to one or more loads behind-the-meter. For instance, renewable energy sources (e.g., wind, solar, hydroelectric, wave, water current, tidal), fossil fuel power generation sources (coal, natural gas), and other types of power producers (e.g., nuclear power) may be positioned in an arrangement that enables the intermittent supply of generated power behind-the-meter to one or more BTM loads. One of ordinary skill in the art will recognize that the generation station **202** may vary based on an application or design in accordance with one or more example embodiments.

In addition, the particular arrangement (e.g., connections) between the generation station and one or more BTM loads can vary within examples. In one embodiment, a generation station may be positioned in an arrangement wherein the generation station selectively supplies power to the grid and/or to one or more BTM loads. As such, power cost-analysis and other factors (e.g., predicted weather condi-

12

tions, contractual obligations, etc.) may be used by the generation station, a BTM load control system, a remote master control system, or some other system or enterprise, to selectively output power to either the grid or to one or more BTM loads in a manner that maximizes revenue to the generation station. In such an arrangement, the generation station may also be able to supply both the grid and one or more BTM loads simultaneously. In some instances, the arrangement may be configured to allow dynamic manipulation of the percentage of the overall generated power that is supplied to each option at a given time. For example, in some time periods, the generation station may supply no power to the BTM load.

In addition, the type of loads that are positioned behind-the-meter can vary within example embodiments. In general, a load that is behind-the-meter may correspond to any type of load capable of receiving and utilizing power behind-the-meter from a generation station. Some examples of loads include, but are not limited to, datacenters and electric vehicle (EV) charging stations.

Preferred BTM loads are loads that can be subject to intermittent power supply because BTM power may be available intermittently. In some instances, the generation station may generate power intermittently. For example, wind power station **102c** and/or photovoltaic power station **102d** may only generate power when resource are available or favorable. Additionally or alternatively, BTM power availability at a generation station may only be available intermittently due to power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and/or operational directives from grid operators or generation station operators.

Some example embodiments of BTM loads described herein involve using one or more computing systems to serve as a BTM load at a generation station. In particular, the computing system or computing systems may receive power behind-the-meter from the generation station to perform various computational operations, such as processing or storing information, performing calculations, mining for cryptocurrencies, supporting blockchain ledgers, and/or executing applications, etc.

Multiple computing systems positioned behind-the-meter may operate as part of a "flexible" datacenter that is configured to operate only intermittently and to receive and utilize BTM power to carry out various computational operations similar to a traditional datacenter. In particular, the flexible datacenter may include computing systems and other components (e.g., support infrastructure, a control system) configured to utilize BTM power from one or more generation stations. The flexible datacenter may be configured to use particular load ramping abilities (e.g., quickly increase or decrease power usage) to effectively operate during intermittent periods of time when power is available from a generation station and supplied to the flexible datacenter behind-the-meter, such as during situations when supplying generated power to the grid is not favorable for the generation station.

In some instances, the amount of power consumed by the computing systems at a flexible datacenter can be ramped up and down quickly, and potentially with high granularity (i.e., the load can be changed in small increments if desired). This may be done based on monitored power system conditions or other information analyses as discussed herein. As recited above, this can enable a generation station to avoid negative power market pricing and to respond quickly to grid direc-

US 10,608,433 B1

13

tives. And by extension, the flexible datacenter may obtain BTM power at a price lower than the cost for power from the grid.

Various types of computing systems can provide granular power ramping. Preferably, the computing systems can perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp down or up. In some embodiments, a control system may be used to activate or de-activate one or more computing systems in an array of computing systems. For example, the control system may provide control instructions to one or more blockchain miners (e.g., a group of blockchain miners), including instructions for powering on or off, adjusting frequency of computing systems performing operations (e.g., adjusting the processing frequency), adjusting the quantity of operations being performed, and when to operate within a low power mode (if available).

Within examples, a control system may correspond to a specialized computing system or may be a computing system within a datacenter serving in the role of the control system. The location of the control system can vary within examples as well. For instance, the control system may be located at a datacenter or physically separate from the datacenter. In some examples, the control system may be part of a network of control systems that manage computational operations, power consumption, and other aspects of a fleet of datacenters. The fleet of datacenters may include one or more traditional datacenters and/or flexible datacenters.

Some embodiments may involve using one or more control systems to direct time-insensitive (e.g., interruptible) computational tasks to computational hardware, such as central processing units (CPUs) and graphics processing units (GPUs), sited behind the meter, while other hardware is sited in front of the meter (i.e., consuming metered grid power via a customer meter (e.g., 106d, 1060) and possibly remote from the behind-the-meter hardware. As such, parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, machine learning, neural network processing, queued operations, and oil and gas field simulation models, are good candidates for such interruptible computational operations.

FIG. 2 shows a behind-the-meter arrangement with optional grid-power, including one or more flexible datacenters, according to one or more example embodiments. Dark arrows illustrate a typical power delivery direction. Consistent with FIG. 1, the arrangement illustrates a generation station 202 in the generation segment 102 of a Wide-Area Synchronous Grid. The generation station 202 supplies utility-scale power (typically >50 MW) via a generation power connection 250 to the Point of Interconnection 103 between the generation station 202 and the rest of the grid. Typically, the power supplied on connection 250 may be at 34.5 kV AC, but it may be higher or lower. Depending on the voltage at connection 250 and the voltage at transmission lines 104a, a transformer system 203 may step up the power supplied from the generation station 202 to high voltage (e.g., 115 kV+AC) for transmission over connection 252 and onto transmission lines 104a of transmission segment 104. Grid power carried on the transmission segment 104 may be from generation station 202 as well as other generation stations (not shown). Also consistent with FIG. 1, grid power is consumed at one or more distribution networks, including example distribution network 206. Grid power may be taken from the transmission lines 104a via connector 254 and stepped down to distribution network

14

voltages (e.g., typically 4 kV to 26 kV AC) and sent into the distribution networks, such as distribution network 206 via distribution line 256. The power on distribution line 256 may be further stepped down (not shown) before entering individual consumer facilities such as a remote master control system 262 and/or traditional datacenters 260 via customer meters 206A, which may correspond to customer meters 106d in FIG. 1, or customer meters 106f in FIG. 1 if the respective consumer facility includes a local customer power system, such as 106e (not shown in FIG. 2).

Consistent with FIG. 1, power entering the grid from generation station 202 is metered by a utility-scale generation-side meter. A utility-scale generation-side meter 253 is shown on the low side of transformer system 203 and an alternative location is shown as 253A on the high side of transformer system 203. Both locations may be considered settlement metering points for the generation station 202 at the POI 103. Alternatively, a utility-scale generation-side meter for the generation station 202 may be located at another location consistent with the descriptions of such meters provided herein.

Generation station 202 includes power generation equipment 210, which may include, as examples, wind turbines and/or photovoltaic panels. Power generation equipment 210 may further include other electrical equipment, including but not limited to switches, busses, collectors, inverters, and power unit transformers (e.g., transformers in wind turbines).

As illustrated in FIG. 2, generation station 202 is configured to connect with BTM equipment which may function as BTM loads. In the illustrated embodiment of FIG. 2, the BTM equipment includes flexible datacenters 220. Various configurations to supply BTM power to flexible datacenters 220 within the arrangement of FIG. 2 are described herein.

In one configuration, generated power may travel from the power generation equipment 210 over one or more connectors 230A, 230B to one or more electrical busses 240A, 240B, respectively. Each of the connectors 230A, 230B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 230B is shown with an open switch, and connector 230A is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used in various embodiments when BTM power is supplied without significant power conversion to BTM loads.

In various configurations, the busses 240A and 240B may be separated by an open switch 240C or combined into a common bus by a closed switch 240C.

In another configuration, generated power may travel from the power generation equipment 210 to the high side of a local step-down transformer 214. The generated power may then travel from the low side of the local step-down transformer 214 over one or more connectors 232A, 232B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 232A, 232B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 232A is shown with an open switch, and connector 232B is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the power generation equipment 210, but the generated power must be stepped down prior to use at the BTM loads.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of a local step-up transformer 212. The generated power may

US 10,608,433 B1

15

then travel from the high side of the local step-up transformer 212 over one or more connectors 234A, 234B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 234A, 234B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 234A, 234B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212.

In another configuration, generated power may travel from the power generation equipment 210 to the low side of the local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 to the high side of local step-down transformer 213. The generated power may then travel from the low side of the local step-down transformer 213 over one or more connectors 236A, 236B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 236A, 236B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 236A, 236B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212, but the power must be stepped down prior to use at the BTM loads.

In one embodiment, power generated at the generation station 202 may be used to power a generation station control system 216 located at the generation station 202, when power is available. The generation station control system 216 may typically control the operation of the generation station 202. Generated power used at the generation station control system 216 may be supplied from bus 240A via connector 216A and/or from bus 240B via connector 216B. Each of the connectors 216A, 216B may be a switched connector such that power may be routed independently to 240A and/or 240B. While the generation station control system 216 can consume BTM power when powered via bus 240A or bus 240B, the BTM power taken by generation station control system 216 is insignificant in terms of rendering an economic benefit. Further, the generation station control system 216 is not configured to operate intermittently, as it generally must remain always on. Further still, the generation station control system 216 does not have the ability to quickly ramp a BTM load up or down.

In another embodiment, grid power may alternatively or additionally be used to power the generation station control system 216. As illustrated here, metered grid power from a distribution network, such as distribution network 206 for simplicity of illustration purposes only, may be used to power generation station control system 216 over connector 216C. Connector 216C may be a switched connector so that metered grid power to the generation station control system 216 can be switched on or off as needed. More commonly, metered grid power would be delivered to the generation station control system 216 via a separate distribution network (not shown), and also over a switched connector. Any such grid power delivered to the generation station control system 216 is metered by a customer meter 206A and subject to T&D costs.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202

16

through POI 103 and such grid power may power the generation station control system 216.

In some configurations, an energy storage system 218 may be connected to the generation station 202 via connector 218A, which may be a switched connector. For illustrative purposes only, connector 218A is shown with an open switch but in some embodiments it may be closed. The energy storage system 218 may be connected to bus 240A and/or bus 240B and store energy produced by the power generation equipment 210. The energy storage system may also be isolated from generation station 202 by switch 242A. In times of need, such as when the power generation equipment in an idle or off state and not generating power, the energy storage system may feed power to, for example, the flexible datacenters 220. The energy storage system may also be isolated from the flexible datacenters 220 by switch 242B.

In a preferred embodiment, as illustrated, power generation equipment 210 supplies BTM power via connector 242 to flexible datacenters 220. The BTM power used by the flexible datacenters 220 was generated by the generation station 202 and did not pass through the POI 103 or utility-scale generation-side meter 253, and is not subject to T&D charges. Power received at the flexible datacenters 220 may be received through respective power input connectors 220A. Each of the respective connectors 220A may be a switched connector that can electrically isolate the respective flexible datacenter 220 from the connector 242. Power equipment 220B may be arranged between the flexible datacenters 220 and the connector 242. The power equipment 220B may include, but is not limited to, power conditioners, unit transformers, inverters, and isolation equipment. As illustrated, each flexible datacenter 220 may be served by a respective power equipment 220B. However, in another embodiment, one power equipment 220B may serve multiple flexible datacenter 220.

In one embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter and electrically connected to the power generation equipment 210 behind (i.e., prior to) the generation station's POI 103 with the rest of the electrical grid.

In one embodiment, BTM power produced by the power generation equipment 210 is utilized by the flexible datacenters 220 behind (i.e., prior to) the generation station's POI with an electrical grid.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as the flexible datacenters 220 are electrically connected to the generation station 202, and generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), and the flexible datacenters 220 receive power from the generation station 202, but the power received by the flexible datacenters 220 from the generation station 202 has not passed through a utility-scale generation-side meter. In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where power produced at the generation station 202 is subject to metering by utility-scale generation-side meter

US 10,608,433 B1

17

253 (or 253A, or another utility-scale generation-side meter), but the BTM power supplied to the flexible datacenters 220 is utilized before being metered at the utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter). In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202—for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as they are electrically connected to the generation station 202 that supplies power to the grid, and the flexible datacenters 220 receive power from the generation station 202 that is not subject to T&D charges, but power otherwise received from the grid that is supplied by the generation station 202 is subject to T&D charges.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 that supplies power to a grid, and the generated power is not subject to T&D charges before being used by flexible datacenters 220, but power otherwise received from the connected grid is subject to T&D charges.

In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter because they receive power generated from the generation station 202 intended for the grid, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 for distribution to the grid, and the flexible datacenters 220 receive that power, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

In another embodiment, metered grid power may alternatively or additionally be used to power one or more of the flexible datacenters 220, or a portion within one or more of the flexible datacenters 220. As illustrated here for simplicity, metered grid power from a distribution network, such as distribution network 206, may be used to power one or more flexible datacenters 220 over connector 256A and/or 256B. Each of connector 256A and/or 256B may be a switched connector so that metered grid power to the flexible datacenters 220 can be switched on or off as needed. More commonly, metered grid power would be delivered to the flexible datacenters 220 via a separate distribution network (not shown), and also over switched connectors. Any such grid power delivered to the flexible datacenters 220 is metered by customer meters 206A and subject to T&D costs. In one embodiment, connector 256B may supply metered grid power to a portion of one or more flexible datacenters 220. For example, connector 256B may supply metered grid power to control and/or communication systems for the flexible datacenters 220 that need constant power and cannot be subject to intermittent BTM power. Connector 242 may supply solely BTM power from the generation station 202 to high power demand computing systems within the flexible datacenters 220, in which case at least a portion of each flexible datacenters 220 so connected is operating as a BTM load. In another embodiment, connector 256A and/or 256B may supply all power used at one or more of the flexible

18

datacenters 220, in which case each of the flexible datacenters 220 so connected would not be operating as a BTM load.

In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the flexible datacenters 220.

The flexible datacenters 220 are shown in an example arrangement relative to the generation station 202. Particularly, generated power from the generation station 202 may be supplied to the flexible datacenters 220 through a series of connectors and/or busses (e.g., 232B, 240B, 242, 220A). As illustrated, in other embodiments, connectors between the power generation equipment 210 and other components may be switched open or closed, allowing other pathways for power transfer between the power generation equipment 210 and components, including the flexible datacenters 220. Additionally, the connector arrangement shown is illustrative only and other circuit arrangements are contemplated within the scope of supplying BTM power to a BTM load at generation station 202. For example, there may be more or fewer transformers, or one or more of transformers 212, 213, 214 may be transformer systems with multiple steppings and/or may include additional power equipment including but not limited to power conditioners, filters, switches, inverters, and/or AC/DC-DC/AC isolators. As another example, metered grid power connections to flexible datacenters 220 are shown via both 256A and 256B; however, a single connection may connect one or more flexible datacenters 220 (or power equipment 220B) to metered grid power and the one or more flexible datacenters 220 (or power equipment 220B) may include switching apparatus to direct BTM power and/or metered grid power to control systems, communication systems, and/or computing systems as desired.

In some examples, BTM power may arrive at the flexible datacenters 220 in a three-phase AC format. As such, power equipment (e.g., power equipment 220B) at one or more of the flexible datacenters 220 may enable each flexible datacenter 220 to use one or more phases of the power. For instance, the flexible datacenters 220 may utilize power equipment (e.g., power equipment 220B, or alternatively or additionally power equipment that is part of the flexible datacenter 220) to convert BTM power received from the generation station 202 for use at computing systems at each flexible datacenter 220. In other examples, the BTM power may arrive at one or more of the flexible datacenters 220 as DC power. As such, the flexible datacenters 220 may use the DC power to power computing systems. In some such examples, the DC power may be routed through a DC-to-DC converter that is part of power equipment 220B and/or flexible datacenter 220.

In some configurations, a flexible datacenter 220 may be arranged to only have access to power received behind-the-meter from a generation station 202. In the arrangement of FIG. 2, the flexible datacenters 220 may be arranged only with a connection to the generation station 202 and depend solely on power received behind-the-meter from the generation station 202. Alternatively or additionally, the flexible datacenters 220 may receive power from energy storage system 218.

In some configurations, one or more of the flexible datacenters 220 can be arranged to have connections to multiple sources that are capable of supplying power to a flexible datacenter 220. To illustrate a first example, the flexible datacenters 220 are shown connected to connector 242, which can be connected or disconnected via switches to

US 10,608,433 B1

19

the energy storage system 218 via connector 218A, the generation station 202 via bus 240B, and grid power via metered connector 256A. In one embodiment, the flexible datacenters 220 may selectively use power received behind-the-meter from the generation station 202, stored power supplied by the energy storage system 218, and/or grid power. For instance, flexible datacenters 220 may use power stored in the energy storage system 218 when costs for using power supplied behind-the-meter from the generation station 202 are disadvantageous. By having access to the energy storage system 218 available, the flexible datacenters 220 may use the stored power and allow the generation station 202 to subsequently refill the energy storage system 218 when cost for power behind-the-meter is low. Alternatively, the flexible datacenters 220 may use power from multiple sources simultaneously to power different components (e.g., a first set and a second set of computing systems). Thus, the flexible datacenters 220 may leverage the multiple connections in a manner that can reduce the cost for power used by the computing systems at the flexible datacenters 220. The flexible datacenters 220 control system or the remote master control system 262 may monitor power conditions and other factors to determine whether the flexible datacenters 220 should use power from either the generation station 202, grid power, the energy storage system 218, none of the sources, or a subset of sources during a given time range. Other arrangements are possible as well. For example, the arrangement of FIG. 2 illustrates each flexible datacenter 220 as connected via a single connector 242 to energy storage system 218, generation station 202, and metered grid power via 256A. However, one or more flexible datacenters 220 may have independent switched connections to each energy source, allowing the one or more flexible datacenters 220 to operate from different energy sources than other flexible datacenters 220 at the same time.

The selection of which power source to use at a flexible datacenter (e.g., the flexible datacenters 220) or another type of BTM load can change based on various factors, such as the cost and availability of power from both sources, the type of computing systems using the power at the flexible datacenters 220 (e.g., some systems may require a reliable source of power for a long period), the nature of the computational operations being performed at the flexible datacenters 220 (e.g., a high priority task may require immediate completion regardless of cost), and temperature and weather conditions, among other possible factors. As such, a datacenter control system at the flexible datacenters 220, the remote master control system 262, or another entity (e.g., an operator at the generation station 202) may also influence and/or determine the source of power that the flexible datacenters 220 use at a given time to complete computational operations.

In some example embodiments, the flexible datacenters 220 may use power from the different sources to serve different purposes. For example, the flexible datacenters 220 may use metered power from grid power to power one or more systems at the flexible datacenters 220 that are configured to be always-on (or almost always on), such as a control and/or communication system and/or one or more computing systems (e.g., a set of computing systems performing highly important computational operations). The flexible datacenters 220 may use BTM power to power other components within the flexible datacenters 220, such as one or more computing systems that perform less critical computational operations.

In some examples, one or more flexible datacenters 220 may be deployed at the generation station 202. In other

20

examples, flexible datacenters 220 may be deployed at a location geographically remote from the generation station 202, while still maintaining a BTM power connection to the generation station 202.

In another example arrangement, the generation station 202 may be connected to a first BTM load (e.g., a flexible datacenter 220) and may supply power to additional BTM loads via connections between the first BTM load and the additional BTM loads (e.g., a connection between a flexible datacenter 220 and another flexible datacenter 220).

The arrangement in FIG. 2, and components included therein, are for non-limiting illustration purposes and other arrangements are contemplated in examples. For instance, in another example embodiment, the arrangement of FIG. 2 may include more or fewer components, such as more BTM loads, different connections between power sources and loads, and/or a different number of datacenters. In addition, some examples may involve one or more components within the arrangement of FIG. 2 being combined or further divided.

Within the arrangement of FIG. 2, a control system, such as the remote master control system 262 or another component (e.g., a control system associated with the grid operator, the generation station control system 216, or a datacenter control system associated with a traditional datacenter or one or more flexible datacenters) may use information to efficiently manage various operations of some of the components within the arrangement of FIG. 2. For example, the remote master control system 262 or another component may manage distribution and execution of computational operations at one or more traditional datacenters 260 and/or flexible datacenters 220 via one or more information-processing algorithms. These algorithms may utilize past and current information in real-time to manage operations of the different components. These algorithms may also make some predictions based on past trends and information analysis. In some examples, multiple computing systems may operate as a network to process information.

Information used to make decisions may include economic and/or power-related information, such as monitored power system conditions. Monitored power system conditions may include one or more of excess power generation at a generation station 202, excess power at a generation station 202 that a connected grid cannot receive, power generation at a generation station 202 subject to economic curtailment, power generation at a generation station 202 subject to reliability curtailment, power generation at a generation station 202 subject to power factor correction, low power generation at a generation station 202, start up conditions at a generation station 202, transient power generation conditions at a generation station 202, or testing conditions where there is an economic advantage to using behind-the-meter power generation at a generation station 202. These different monitored power system conditions can be weighted differently during processing and analysis.

In some examples, the information can include the cost for power from available sources (e.g., BTM power at the generation station 202 versus metered grid power) to enable comparisons to be made which power source costs less. In some instances, the information may include historic prices for power to enable the remote master control system 262 or another system to predict potential future prices in similar situations (e.g., the cost of power tends to trend upwards for grid power during warmer weather and peak-use hours). The information may also indicate the availability of power from the various sources (e.g., BTM power at the generation

US 10,608,433 B1

21

station 262, the energy storage system 218 at the generation station 262, and/or metered grid power).

In addition, the information may also include other data, including information associated with operations at components within the arrangement. For instance, the information may include data associated with performance of operations at the flexible datacenters 220 and the traditional datacenters 260, such as the number of computational tasks currently being performed, the types of tasks being performed (e.g., type of computational operation, time-sensitivity, etc.), the number, types, and capabilities of available computing systems, the amount of computational tasks awaiting performance, and the types of computing systems at one or more datacenters, among others. The information may also include data specifying the conditions at one or more datacenters (e.g., whether or not the temperatures are in a desired range, the amount of power available within an energy storage system such as 218), the amount of computational tasks awaiting performance in the queue of one or more of the datacenters, and the identities of the entities associated with the computational operations at one or more of the datacenters. Entities associated with computational operations may be, for example, owners of the datacenters, customers who purchase computational time at the datacenters, or other entities.

The information used by the remote master control system 262 or another component may include data associated with the computational operations to be performed, such as deadlines, priorities (e.g., high vs. low priority tasks), cost to perform based on required computing systems, the optimal computing systems (e.g., CPU vs GPU vs ASIC; processing unit capabilities, speeds, or frequencies, or instructional sets executable by the processing units) for performing each requested computational task, and prices each entity (e.g., company) is willing to pay for computational operations to be performed or otherwise supported via computing systems at a traditional datacenter 260 or a flexible datacenter 220, among others. In addition, the information may also include other data (e.g., weather conditions at locations of datacenters or power sources, any emergencies associated with a datacenter or power source, or the current value of bids associated with an auction for computational tasks).

The information may be updated in-real time and used to make the different operational decisions within the arrangement of FIG. 2. For instance, the information may help a component (e.g., the remote master control system 262 or a control system at a flexible datacenter 220) determine when to ramp up or ramp down power use at a flexible datacenter 220 or when to switch one or more computing systems at a flexible datacenter 220 into a low power mode or to operate at a different frequency, among other operational adjustments. The information can additionally or alternatively help a component within the arrangement of FIG. 2 to determine when to transfer computational operations between computing systems or between datacenters based on various factors. In some instances, the information may also be used to determine when to temporarily stop performing a computational operation or when to perform a computational operation at multiple sites for redundancy or other reasons. The information may further be used to determine when to accept new computational operations from entities or when to temporarily suspend accepting new tasks to be performed due to lack of computing system availability.

The remote master control system 262 represents a computing system that is capable of obtaining, managing, and using the information described above to manage and over-

22

see one or more operations within the arrangement of FIG. 2. As such, the remote master control system 262 may be one or more computing systems configured to process all, or a subset of, the information described above, such as power, environment, computational characterization, and economic factors to assist with the distribution and execution of computing operations among one or more datacenters. For instance, the remote master control system 262 may be configured to obtain and delegate computational operations among one or more datacenters based on a weighted analysis of a variety of factors, including one or more of the cost and availability of power, the types and availability of the computing systems at each datacenter, current and predicted weather conditions at the different locations of flexible datacenters (e.g., flexible datacenters 220) and generation stations (e.g., generation stations 202), levels of power storage available at one or more energy storage systems (e.g., energy storage system 218), and deadlines and other attributes associated with particular computational operations, among other possible factors. As such, the analysis of information performed by the remote master control system 262 may vary within examples. For instance, the remote master control system 262 may use real-time information to determine whether or not to route a computational operation to a particular flexible datacenter (e.g., a flexible datacenter 220) or to transition a computational operation between datacenters (e.g., from traditional datacenter 260 to a flexible datacenter 220).

As shown in FIG. 2, the generation station 202 may be able to supply power to the grid and/or BTM loads such as flexible datacenters 220. With such a configuration, the generation station 202 may selectively provide power to the BTM loads and/or the grid based on economic and power availability considerations. For example, the generation station 202 may supply power to the grid when the price paid for the power exceeds a particular threshold (e.g., the power price offered by operators of the flexible datacenters 220). In some instances, the operator of a flexible datacenter and the operator of a generation station capable of supplying BTM power to the flexible datacenter may utilize a predefined arrangement (e.g., a contract) that specifies a duration and/or price range when the generation station may supply power to the flexible datacenter.

The remote master control system 262 may be capable of directing one or more flexible datacenters 220 to ramp-up or ramp-down to desired power consumption levels, and/or to control cooperative action of multiple flexible datacenters by determining how to power each individual flexible datacenter 220 in accordance with operational directives.

The configuration of the remote master control system 262 can vary within examples as further discussed with respect to FIGS. 2, 3, and 7-9. The remote master control system 262 may operate as a single computing system or may involve a network of computing systems. Preferably, the remote master control system 262 is implemented across one or more servers in a fault-tolerant operating environment that ensures continuous uptime and connectivity by virtue of its distributed nature. Alternatively, although the remote master control system 262 is shown as a physically separate component arrangement for FIG. 2, the remote master control system 262 may be combined with another component in other embodiments. To illustrate an example, the remote master control system 262 may operate as part of a flexible datacenter (e.g., a computing system or a datacenter control system of the flexible datacenter 220), includ-

US 10,608,433 B1

23

ing sharing components with a flexible datacenter, sharing power with a flexible datacenter, and/or being co-located with a flexible datacenter.

In addition, the remote master control system 262 may communicate with components within the arrangement of FIG. 2 using various communication technologies, including wired and wireless communication technologies. For instance, the remote master control system 262 may use wired (not illustrated) or wireless communication to communicate with datacenter control systems or other computing systems at the flexible datacenters 220 and the traditional datacenters 260. The remote master control system 262 may also communicate with entities inside or outside the arrangement of FIG. 2 and other components within the arrangement of FIG. 2 via wired or wireless communication. For instance, the remote master control system 262 may use wireless communication to obtain computational operations from entities seeking support for the computational operations at one or more datacenters in exchange for payment. The remote master control system 262 may communicate directly with the entities or may obtain the computational operations from the traditional datacenters 260. For instance, an entity may submit jobs (e.g., computational operations) to one or more traditional datacenters 260. The remote master control system 262 may determine that transferring one or more of the computational operations to a flexible datacenter 220 may better support the transferred computational operations. For example, the remote master control system 262 may determine that the transfer may enable the computational operations to be completed quicker and/or at a lower cost. In some examples, the remote master control system 262 may communicate with the entity to obtain approval prior to transferring the one or more computational operations.

The remote master control system 262 may also communicate with grid operators and/or an operator of generation station 202 to help determine power management strategies when distributing computational operations across the various datacenters. In addition, the remote master control system 262 may communicate with other sources, such as weather prediction systems, historical and current power price databases, and auction systems, etc.

In further examples, the remote master control system 262 or another computing system within the arrangement of FIG. 2 may use wired or wireless communication to submit bids within an auction that involves a bidder (e.g., the highest bid) obtaining computational operations or other tasks to be performed. Particularly, the remote master control system 262 may use the information discussed above to develop bids to obtain computing operations for performance at available computing systems at flexible datacenters (e.g., flexible datacenters 220).

In the example arrangement shown in FIG. 2, the flexible datacenters 220 represent example loads that can receive power behind-the-meter from the generation station 202. In such a configuration, the flexible datacenters 220 may obtain and utilize power behind-the-meter from the generation station 202 to perform various computational operations. Performance of a computational operation may involve one or more computing systems providing resources useful in the computational operation. For instance, the flexible datacenters 220 may include one or more computing systems configured to store information, perform calculations and/or parallel processes, perform simulations, mine cryptocurrencies, and execute applications, among other potential tasks. The computing systems can be specialized or generic and can be arranged at each flexible datacenter 220 in a variety

24

of ways (e.g., straight configuration, zig-zag configuration) as further discussed with respect to FIGS. 6A, 6B. Furthermore, although the example arrangement illustrated in FIG. 2 shows configurations where flexible datacenters 220 serve as BTM loads, other types of loads can be used as BTM loads within examples.

The arrangement of FIG. 2 includes the traditional datacenters 260 coupled to metered grid power. The traditional datacenters 260 using metered grid power to provide computational resources to support computational operations. One or more enterprises may assign computational operations to the traditional datacenters 260 with expectations that the datacenters reliably provide resources without interruption (i.e., non-intermittently) to support the computational operations, such as processing abilities, networking, and/or volatile storage. Similarly, one or more enterprises may also request computational operations to be performed by the flexible datacenters 220. The flexible datacenters 220 differ from the traditional datacenters 260 in that the flexible datacenters 220 are arranged and/or configured to be connected to BTM power, are expected to operate intermittently, and are expected to ramp load (and thus computational capability) up or down regularly in response to control directives. In some examples, the flexible datacenters 220 and the traditional datacenters 260 may have similar configurations and may only differ based on the source(s) of power relied upon to power internal computing systems. Preferably, however, the flexible datacenters 220 include particular fast load ramping abilities (e.g., quickly increase or decrease power usage) and are intended and designed to effectively operate during intermittent periods of time.

FIG. 3 shows a block diagram of the remote master control system 300 according to one or more example embodiments. Remote master control system 262 may take the form of remote master control system 300, or may include less than all components in remote master control system 300, different components than in remote master control system 300, and/or more components than in remote master control system 300.

The remote master control system 300 may perform one or more operations described herein and may include a processor 302, a data storage unit 304, a communication interface 306, a user interface 308, an operations and environment analysis module 310, and a queue system 312. In other examples, the remote master control system 300 may include more or fewer components in other possible arrangements.

As shown in FIG. 3, the various components of the remote master control system 300 can be connected via one or more connection mechanisms (e.g., a connection mechanism 314). In this disclosure, the term "connection mechanism" means a mechanism that facilitates communication between two or more devices, systems, components, or other entities. For instance, a connection mechanism can be a simple mechanism, such as a cable, PCB trace, or system bus, or a relatively complex mechanism, such as a packet-based communication network (e.g., LAN, WAN, and/or the Internet). In some instances, a connection mechanism can include a non-tangible medium (e.g., where the connection is wireless).

As part of the arrangement of FIG. 2, the remote master control system 300 (corresponding to remote master control system 262) may perform a variety of operations, such as management and distribution of computational operations among datacenters, monitoring operational, economic, and environment conditions, and power management. For instance, the remote master control system 300 may obtain

US 10,608,433 B1

25

computational operations from one or more enterprises for performance at one or more datacenters. The remote master control system 300 may subsequently use information to distribute and assign the computational operations to one or more datacenters (e.g., the flexible datacenters 220) that have the resources (e.g., particular types of computing systems and available power) available to complete the computational operations. In some examples, the remote master control system 300 may assign all incoming computational operation requests to the queue system 312 and subsequently assign the queued requests to computing systems based on an analysis of current market and power conditions.

Although the remote master control system 300 is shown as a single entity, a network of computing systems may perform the operations of the remote master control system 300 in some examples. For example, the remote master control system 300 may exist in the form of computing systems (e.g., datacenter control systems) distributed across multiple datacenters.

The remote master control system 300 may include one or more processors 302. As such, the processor 302 may represent one or more general-purpose processors (e.g., a microprocessor) and/or one or more special-purpose processors (e.g., a digital signal processor (DSP)). In some examples, the processor 302 may include a combination of processors within examples. The processor 302 may perform operations, including processing data received from the other components within the arrangement of FIG. 2 and data obtained from external sources, including information such as weather forecasting systems, power market price systems, and other types of sources or databases.

The data storage unit 304 may include one or more volatile, non-volatile, removable, and/or non-removable storage components, such as magnetic, optical, or flash storage, and/or can be integrated in whole or in part with the processor 302. As such, the data storage unit 304 may take the form of a non-transitory computer-readable storage medium, having stored thereon program instructions (e.g., compiled or non-compiled program logic and/or machine code) that, when executed by the processor 302, cause the remote master control system 300 to perform one or more acts and/or functions, such as those described in this disclosure. Such program instructions can define and/or be part of a discrete software application. In some instances, the remote master control system 300 can execute program instructions in response to receiving an input, such as from the communication interface 306, the user interface 308, or the operations and environment analysis module 310. The data storage unit 304 may also store other information, such as those types described in this disclosure.

In some examples, the data storage unit 304 may serve as storage for information obtained from one or more external sources. For example, data storage unit 304 may store information obtained from one or more of the traditional datacenters 260, a generation station 202, a system associated with the grid, and flexible datacenters 220. As examples only, data storage 304 may include, in whole or in part, local storage, dedicated server-managed storage, network attached storage, and/or cloud-based storage, and/or combinations thereof.

The communication interface 306 can allow the remote master control system 300 to connect to and/or communicate with another component according to one or more protocols. For instance, the communication interface 306 may be used to obtain information related to current, future, and past prices for power, power availability, current and predicted

26

weather conditions, and information regarding the different datacenters (e.g., current workloads at datacenters, types of computing systems available within datacenters, price to obtain power at each datacenter, levels of power storage available and accessible at each datacenter, etc.). In an example, the communication interface 306 can include a wired interface, such as an Ethernet interface or a high-definition serial-digital-interface (HD-SDI). In another example, the communication interface 406 can include a wireless interface, such as a cellular, satellite, WiMAX, or WI-FI interface. A connection can be a direct connection or an indirect connection, the latter being a connection that passes through and/or traverses one or more components, such as such as a router, switcher, or other network device. Likewise, a wireless transmission can be a direct transmission or an indirect transmission. The communication interface 306 may also utilize other types of wireless communication to enable communication with datacenters positioned at various locations.

The communication interface 306 may enable the remote master control system 300 to communicate with the components of the arrangement of FIG. 2. In addition, the communication interface 306 may also be used to communicate with the various datacenters, power sources, and different enterprises submitting computational operations for the datacenters to support.

The user interface 308 can facilitate interaction between the remote master control system 300 and an administrator or user, if applicable. As such, the user interface 308 can include input components such as a keyboard, a keypad, a mouse, a touch-sensitive panel, a microphone, and/or a camera, and/or output components such as a display device (which, for example, can be combined with a touch-sensitive panel), a sound speaker, and/or a haptic feedback system. More generally, the user interface 308 can include hardware and/or software components that facilitate interaction between remote master control system 300 and the user of the system.

In some examples, the user interface 308 may enable the manual examination and/or manipulation of components within the arrangement of FIG. 2. For instance, an administrator or user may use the user interface 308 to check the status of, or change, one or more computational operations, the performance or power consumption at one or more datacenters, the number of tasks remaining within the queue system 312, and other operations. As such, the user interface 308 may provide remote connectivity to one or more systems within the arrangement of FIG. 2.

The operations and environment analysis module 310 represents a component of the remote master control system 300 associated with obtaining and analyzing information to develop instructions/directives for components within the arrangement of FIG. 2. The information analyzed by the operations and environment analysis module 310 can vary within examples and may include the information described above with respect predicting and/or directing the use of BTM power. For instance, the operations and environment analysis module 310 may obtain and access information related to the current power state of computing systems operating as part of the flexible datacenters 220 and other datacenters that the remote master control system 300 has access to. This information may be used to determine when to adjust power usage or mode of one or more computing systems. In addition, the remote master control system 300 may provide instructions a flexible datacenter 220 to cause a subset of the computing systems to transition into a low power mode to consume less power while still performing

operations at a slower rate. The remote master control system 300 may also use power state information to cause a set of computing systems at a flexible datacenter 220 to operate at a higher power consumption mode. In addition, the remote master control system 300 may transition computing systems into sleep states or power on/off based on information analyzed by the operations and environment analysis module 310.

In some examples, the operations and environment analysis module 310 may use location, weather, activity levels at the flexible datacenters or the generation station, and power cost information to determine control strategies for one or more components in the arrangement of FIG. 2. For instance, the remote master control system 300 may use location information for one or more datacenters to anticipate potential weather conditions that could impact access to power. In addition, the operations and environment analysis module 310 may assist the remote master control system 300 determine whether to transfer computational operations between datacenters based on various economic and power factors.

The queue system 312 represents a queue capable of organizing computational operations to be performed by one or more datacenters. Upon receiving a request to perform a computational operation, the remote master control system 300 may assign the computational operation to the queue until one or more computing systems are available to support the computational operation. The queue system 312 may be used for organizing and transferring computational tasks in real time.

The organizational design of the queue system 312 may vary within examples. In some examples, the queue system 312 may organize indications (e.g., tags, pointers) to sets of computational operations requested by various enterprises. The queue system 312 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 312 may include one or more queues that operate using the FIFO data structure.

In some examples, one or more queues within the queue system 312 may use other designs of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others. Examples using a queue system are further described below with respect to FIG. 9.

In some examples, the remote master control system 300 may be configured to monitor one or more auctions to obtain computational operations for datacenters to support. Particularly, the remote master control system 300 may use resource availability and power prices to develop and submit bids to an external or internal auction system for the right to support particular computational operations. As a result, the remote master control system 300 may identify computational operations that could be supported at one or more flexible datacenters 220 at low costs.

FIG. 4 is a block diagram of a generation station 400, according to one or more example embodiments. Generation station 202 may take the form of generation station 400, or may include less than all components in generation station 400, different components than in generation station 400, and/or more components than in generation station 400. The generation station 400 includes power generation equipment

401, a communication interface 408, a behind-the-meter interface 406, a grid interface 404, a user interface 410, a generation station control system 414, and power transformation equipment 402. The power generation equipment 210 may take the form of power generation equipment 401, or may include less than all components in power generation equipment 401, different components than in power generation equipment 401, and/or more components than in power generation equipment 401. Generation station control system 216 may take the form of generation station control system 414, or may include less than all components in generation station control system 414, different components than in generation station control system 414, and/or more components than in generation station control system 414. Some or all of the components generation station 400 may be connected via a communication interface 516. These components are illustrated in FIG. 4 to convey an example configuration for the generation station 400 (corresponding to generation station 202 shown in FIG. 2). In other examples, the generation station 400 may include more or fewer components in other arrangements.

The generation station 400 can correspond to any type of grid-connected utility-scale power producer capable of supplying power to one or more loads. The size, amount of power generated, and other characteristics of the generation station 400 may differ within examples. For instance, the generation station 400 may be a power producer that provides power intermittently. The power generation may depend on monitored power conditions, such as weather at the location of the generation station 400 and other possible conditions. As such, the generation station 400 may be a temporary arrangement, or a permanent facility, configured to supply power. The generation station 400 may supply BTM power to one or more loads and supply metered power to the electrical grid. Particularly, the generation station 400 may supply power to the grid as shown in the arrangement of FIG. 2.

The power generation equipment 401 represents the component or components configured to generate utility-scale power. As such, the power generation equipment 401 may depend on the type of facility that the generation station 400 corresponds to. For instance, the power generation equipment 401 may correspond to electric generators that transform kinetic energy into electricity. The power generation equipment 401 may use electromagnetic induction to generate power. In other examples, the power generation equipment 401 may utilize electrochemistry to transform chemical energy into power. The power generation equipment 401 may use the photovoltaic effect to transform light into electrical energy. In some examples, the power generation equipment 401 may use turbines to generate power. The turbines may be driven by, for example, wind, water, steam or burning gas. Other examples of power production are possible.

The communication interface 408 can enable the generation station 400 to communicate with other components within the arrangement of FIG. 2. As such, the communication interface 408 may operate similarly to the communication interface 306 of the remote master control system 300 and the communication interface 503 of the flexible datacenter 500.

The generation station control system 414 may be one or more computing systems configured to control various aspects of the generation station 400.

The BTM interface 406 is a module configured to enable the power generation equipment 401 to supply BTM power to one or more loads and may include multiple components.

The arrangement of the BTM interface 406 may differ within examples based on various factors, such as the number of flexible datacenters 220 (or 500) coupled to the generation station 400, the proximity of the flexible datacenters 220 (or 500), and the type of generation station 400, among others. In some examples, the BTM interface 406 may be configured to enable power delivery to one or more flexible datacenters positioned near the generation station 400. Alternatively, the BTM interface 406 may also be configured to enable power delivery to one or more flexible datacenters 220 (or 500) positioned remotely from the generation station 400.

The grid interface 404 is a module configured to enable the power generation equipment 401 to supply power to the grid and may include multiple components. As such, the grid interface 404 may couple to one or more transmission lines (e.g., transmission lines 404a shown in FIG. 2) to enable delivery of power to the grid.

The user interface 410 represents an interface that enables administrators and/or other entities to communicate with the generation station 400. As such, the user interface 410 may have a configuration that resembles the configuration of the user interface 308 shown in FIG. 3. An operator may utilize the user interface 410 to control or monitor operations at the generation station 400.

The power transformation equipment 402 represents equipment that can be utilized to enable power delivery from the power generation equipment 401 to the loads and to transmission lines linked to the grid. Example power transformation equipment 402 includes, but is not limited to, transformers, inverters, phase converters, and power conditioners.

FIG. 5 shows a block diagram of a flexible datacenter 500, according to one or more example embodiments. Flexible datacenters 220 may take the form of flexible datacenter 500, or may include less than all components in flexible datacenter 500, different components than in flexible datacenter 500, and/or more components than in flexible datacenter 500. In the example embodiment shown in FIG. 5, the flexible datacenter 500 includes a power input system 502, a communication interface 503, a datacenter control system 504, a power distribution system 506, a climate control system 508, one or more sets of computing systems 512, and a queue system 514. These components are shown connected by a communication bus 528. In other embodiments, the configuration of flexible datacenter 500 can differ, including more or fewer components. In addition, the components within flexible datacenter 500 may be combined or further divided into additional components within other embodiments.

The example configuration shown in FIG. 5 represents one possible configuration for a flexible datacenter. As such, each flexible datacenter may have a different configuration when implemented based on a variety of factors that may influence its design, such as location and temperature that the location, particular uses for the flexible datacenter, source of power supplying computing systems within the flexible datacenter, design influence from an entity (or entities) that implements the flexible datacenter, and space available for the flexible datacenter. Thus, the embodiment of flexible datacenter 220 shown in FIG. 2 represents one possible configuration for a flexible datacenter out of many other possible configurations.

The flexible datacenter 500 may include a design that allows for temporary and/or rapid deployment, setup, and start time for supporting computational operations. For instance, the flexible datacenter 500 may be rapidly

deployed at a location near a source of generation station power (e.g., near a wind farm or solar farm). Rapid deployment may involve positioning the flexible datacenter 500 at a target location and installing and/or configuring one or more racks of computing systems within. The racks may include wheels to enable swift movement of the computing systems. Although the flexible datacenter 500 could theoretically be placed anywhere, transmission losses may be minimized by locating it proximate to BTM power generation.

The physical construction and layout of the flexible datacenter 500 can vary. In some instances, the flexible datacenter 500 may utilize a metal container (e.g., a metal container 602 shown in FIG. 6A). In general, the flexible datacenter 500 may utilize some form of secure weather-proof housing designed to protect interior components from wind, weather, and intrusion. The physical construction and layout of example flexible datacenters are further described with respect to FIGS. 6A-6B.

Within the flexible datacenter 500, various internal components enable the flexible datacenter 500 to utilize power to perform some form of operations. The power input system 502 is a module of the flexible datacenter 500 configured to receive external power and input the power to the different components via assistance from the power distribution system 506. As discussed with respect to FIG. 2, the sources of external power feeding a flexible datacenter can vary in both quantity and type (e.g., the generation stations 202, 400, grid-power, energy storage systems). Power input system 502 includes a BTM power input sub-system 522, and may additionally include other power input sub-systems (e.g., a grid-power input sub-system 524 and/or an energy storage input sub-system 526). In some instances, the quantity of power input sub-systems may depend on the size of the flexible datacenter and the number and/or type of computing systems being powered. In an example embodiment, the flexible datacenter may use grid power as the primary power supply.

In some embodiments, the power input system 502 may include some or all of flexible datacenter Power Equipment 220B. The power input system 502 may be designed to obtain power in different forms (e.g., single phase or three-phase behind-the-meter alternating current ("AC") voltage, and/or direct current ("DC") voltage). As shown, the power input system 502 includes a BTM power input sub-system 522, a grid power input sub-system 524, and an energy input sub-system 526. These sub-systems are included to illustrate example power input sub-systems that the flexible datacenter 500 may utilize, but other examples are possible. In addition, in some instances, these sub-systems may be used simultaneously to supply power to components of the flexible datacenter 500. The sub-systems may also be used based on available power sources.

In some implementations, the BTM power input sub-system 522 may include one or more AC-to-AC step-down transformers used to step down supplied medium-voltage AC to low voltage AC (e.g., 120V to 600V nominal) used to power computing systems 512 and/or other components of flexible datacenter 500. The power input system 502 may also directly receive single-phase low voltage AC from a generation station as BTM power, from grid power, or from a stored energy system such as energy storage system 218. In some implementations, the power input system 502 may provide single-phase AC voltage to the datacenter control system 504 (and/or other components of flexible datacenter 500) independent of power supplied to computing systems 512 to enable the datacenter control system 504 to perform

US 10,608,433 B1

31

management operations for the flexible datacenter 500. For instance, the grid power input sub-system 524 may use grid power to supply power to the datacenter control system 504 to ensure that the datacenter control system 504 can perform control operations and communicate with the remote master control system 300 (or 262) during situations when BTM power is not available. As such, the datacenter control system 504 may utilize power received from the power input system 502 to remain powered to control the operation of flexible datacenter 500, even if the computational operations performed by the computing system 512 are powered intermittently. In some instances, the datacenter control system 504 may switch into a lower power mode to utilize less power while still maintaining the ability to perform some functions.

The power distribution system 506 may distribute incoming power to the various components of the flexible datacenter 500. For instance, the power distribution system 506 may direct power (e.g., single-phase or three-phase AC) to one or more components within flexible datacenter 500. In some embodiments, the power distribution system 506 may include some or all of flexible datacenter Power Equipment 220B.

In some examples, the power input system 502 may provide three phases of three-phase AC voltage to the power distribution system 506. The power distribution system 506 may controllably provide a single phase of AC voltage to each computing system or groups of computing systems 512 disposed within the flexible datacenter 500. The datacenter control system 504 may controllably select which phase of three-phase nominal AC voltage that power distribution system 506 provides to each computing system 512 or groups of computing systems 512. This is one example manner in which the datacenter control system 504 may modulate power delivery (and load at the flexible datacenter 500) by ramping-up flexible datacenter 500 to fully operational status, ramping-down flexible datacenter 500 to offline status (where only datacenter control system 504 remains powered), reducing load by withdrawing power delivery from, or reducing power to, one or more of the computing systems 512 or groups of the computing systems 512, or modulating power factor correction for the generation station 300 (or 202) by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more of the computing systems 512 or groups of the computing systems 512. The datacenter control system 504 may direct power to certain sets of computing systems based on computational operations waiting for computational resources within the queue system 514. In some embodiments, the flexible datacenter 500 may receive BTM DC power to power the computing systems 512.

One of ordinary skill in the art will recognize that a voltage level of three-phase AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase AC voltage, single-phase AC voltage, and DC voltage may vary based on the application or design in accordance with one or more embodiments.

As discussed above, the datacenter control system 504 may perform operations described herein, such as dynamically modulating power delivery to one or more of the computing systems 512 disposed within flexible datacenter 500. For instance, the datacenter control system 504 may modulate power delivery to one or more of the computing

32

systems 512 based on various factors, such as BTM power availability or an operational directive from a generation station 262 or 300 control system, a remote master control system 262 or 300, or a grid operator. In some examples, the datacenter control system 504 may provide computational operations to sets of computing systems 512 and modulate power delivery based on priorities assigned to the computational operations. For instance, an important computational operation (e.g., based on a deadline for execution and/or price paid by an entity) may be assigned to a particular computing system or set of computing systems 512 that has the capacity, computational abilities to support the computational operation. In addition, the datacenter control system 504 may also prioritize power delivery to the computing system or set of computing systems 512.

In some example, the datacenter control system 504 may further provide directives to one or more computing systems to change operations in some manner. For instance, the datacenter control system 504 may cause one or more computing systems 512 to operate at a lower or higher frequency, change clock cycles, or operate in a different power consumption mode (e.g., a low power mode). These abilities may vary depending on types of computing systems 512 available at the flexible datacenter 500. As a result, the datacenter control system 504 may be configured to analyze the computing systems 512 available either on a periodic basis (e.g., during initial set up of the flexible datacenter 500) or in another manner (e.g., when a new computational operation is assigned to the flexible datacenter 500).

The datacenter control system 504 may also implement directives received from the remote master control system 262 or 300. For instance, the remote master control system 262 or 300 may direct the flexible datacenter 500 to switch into a low power mode. As a result, one or more of the computing systems 512 and other components may switch to the low power mode in response.

The datacenter control system 504 may utilize the communication interface 503 to communicate with the remote master control system 262 or 300, other datacenter control systems of other datacenters, and other entities. As such, the communication interface 503 may include components and operate similar to the communication interface 306 of the remote master control system 300 described with respect to FIG. 4.

The flexible datacenter 500 may also include a climate control system 508 to maintain computing systems 512 within a desired operational temperature range. The climate control system 508 may include various components, such as one or more air intake components, an evaporative cooling system, one or more fans, an immersive cooling system, an air conditioning or refrigerant cooling system, and one or more air outtake components. One of ordinary skill in the art will recognize that any suitable heat extraction system configured to maintain the operation of computing systems 512 within the desired operational temperature range may be used.

The flexible datacenter 500 may further include an energy storage system 510. The energy storage system 510 may store energy for subsequent use by computing systems 512 and other components of flexible datacenter 500. For instance, the energy storage system 510 may include a battery system. The battery system may be configured to convert AC voltage to DC voltage and store power in one or more storage cells. In some instances, the battery system may include a DC-to-AC inverter configured to convert DC

US 10,608,433 B1

33

voltage to AC voltage, and may further include an AC phase-converter, to provide AC voltage for use by flexible datacenter 500.

The energy storage system 510 may be configured to serve as a backup source of power for the flexible datacenter 500. For instance, the energy storage system 510 may receive and retain power from a BTM power source at a low cost (or no cost at all). This low-cost power can then be used by the flexible datacenter 500 at a subsequent point, such as when BTM power costs more. Similarly, the energy storage system 510 may also store energy from other sources (e.g., grid power). As such, the energy storage system 510 may be configured to use one or more of the sub-systems of the power input system 502.

In some examples, the energy storage system 510 may be external to the flexible datacenter 500. For instance, the energy storage system 510 may be an external source that multiple flexible datacenters utilize for back-up power.

The computing systems 512 represent various types of computing systems configured to perform computational operations. Performance of computational operations include a variety of tasks that one or more computing systems may perform, such as data storage, calculations, application processing, parallel processing, data manipulation, cryptocurrency mining, and maintenance of a distributed ledger, among others. As shown in FIG. 5, the computing systems 512 may include one or more CPUs 516, one or more GPUs 518, and/or one or more Application-Specific Integrated Circuits (ASIC's) 520. Each type of computing system 512 may be configured to perform particular operations or types of operations.

Due to different performance features and abilities associated with the different types of computing systems, the datacenter control system 504 may determine, maintain, and/or relay this information about the types and/or abilities of the computing systems, quantity of each type, and availability to the remote master control system 262 or 300 on a routine basis (e.g., periodically or on-demand). This way, the remote master control system 262 or 300 may have current information about the abilities of the computing systems 512 when distributing computational operations for performance at one or more flexible datacenters. Particularly, the remote master control system 262 or 300 may assign computational operations based on various factors, such as the types of computing systems available and the type of computing systems required by each computing operation, the availability of the computing systems, whether computing systems can operate in a low power mode, and/or power consumption and/or costs associated with operating the computing systems, among others.

The quantity and arrangement of these computing systems 512 may vary within examples. In some examples, the configuration and quantity of computing systems 512 may depend on various factors, such as the computational tasks that are performed by the flexible datacenter 500. In other examples, the computing systems 512 may include other types of computing systems as well, such as DSPs, SIMDs, neural processors, and/or quantum processors.

As indicated above, the computing systems 512 can perform various computational operations, including in different configurations. For instance, each computing system may perform a particular computational operation unrelated to the operations performed at other computing systems. Groups of the computing systems 512 may also be used to work together to perform computational operations.

In some examples, multiple computing systems may perform the same computational operation in a redundant

34

configuration. This redundant configuration creates a backup that prevents losing progress on the computational operation in situations of a computing failure or intermittent operation of one or more computing systems. In addition, the computing systems 512 may also perform computational operations using a check point system. The check point system may enable a first computing system to perform operations up to a certain point (e.g., a checkpoint) and switch to a second computing system to continue performing the operations from that certain point. The check point system may also enable the datacenter control system 504 to communicate statuses of computational operations to the remote master control system 262 or 300. This can further enable the remote master control system 262 or 300 to transfer computational operations between different flexible datacenters allowing computing systems at the different flexible datacenters to resume support of computational operations based on the check points.

The queue system 514 may operate similar to the queue system 312 of the remote master control system 300 shown in FIG. 3. Particularly, the queue system 514 may help store and organize computational tasks assigned for performance at the flexible datacenter 500. In some examples, the queue system 514 may be part of a distributed queue system such that each flexible datacenter in a fleet of flexible datacenter includes a queue, and each queue system 514 may be able to communicate with other queue systems. In addition, the remote master control system 262 or 300 may be configured to assign computational tasks to the queues located at each flexible datacenter (e.g., the queue system 514 of the flexible datacenter 500). As such, communication between the remote master control system 262 or 300 and the datacenter control system 504 and/or the queue system 514 may allow organization of computational operations for the flexible datacenter 500 to support.

FIG. 6A shows another structural arrangement for a flexible datacenter, according to one or more example embodiments. The particular structural arrangement shown in FIG. 6A may be implemented at flexible datacenter 500. The illustration depicts the flexible datacenter 500 as a mobile container 702 equipped with the power input system 502, the power distribution system 506, the climate control system 508, the datacenter control system 504, and the computing systems 512 arranged on one or more racks 604. These components of flexible datacenter 500 may be arranged and organized according to an example structural region arrangement. As such, the example illustration represents one possible configuration for the flexible datacenter 500, but others are possible within examples.

As discussed above, the structural arrangement of the flexible datacenter 500 may depend on various factors, such as the ability to maintain temperature within the mobile container 602 within a desired temperature range. The desired temperature range may depend on the geographical location of the mobile container 602 and the type and quantity of the computing systems 512 operating within the flexible datacenter 500 as well as other possible factors. As such, the different design elements of the mobile container 602 including the inner contents and positioning of components may depend on factors that aim to maximize the use of space within mobile container 602, lower the amount of power required to cool the computing systems 512, and make setup of the flexible datacenter 500 efficient. For instance, a first flexible datacenter positioned in a cooler geographic region may include less cooling equipment than a second flexible datacenter positioned in a warmer geographic region.

US 10,608,433 B1

35

As shown in FIG. 6A, the mobile container 602 may be a storage trailer disposed on permanent or removable wheels and configured for rapid deployment. In other embodiments, the mobile container 602 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, the mobile container 602 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile flexible datacenter. As such, the flexible datacenter 500 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. And in still other embodiments, the flexible datacenter 500 might not include a mobile container. For example, the flexible datacenter 500 may be situated within a building or another type of stationary environment.

FIG. 6B shows the computing systems 512 in a straight-line configuration for installation within the flexible datacenter 500, according to one or more example embodiments. As indicated above, the flexible datacenter 500 may include a plurality of racks 604, each of which may include one or more computing systems 512 disposed therein. As discussed above, the power input system 502 may provide three phases of AC voltage to the power distribution system 506. In some examples, the power distribution system 506 may controllably provide a single phase of AC voltage to each computing system 512 or group of computing systems 512 disposed within the flexible datacenter 500. As shown in FIG. 6B, for purposes of illustration only, eighteen total racks 604 are divided into a first group of six racks 606, a second group of six racks 608, and a third group of six racks 610, where each rack contains eighteen computing systems 512. The power distribution system (506 of FIG. 5) may, for example, provide a first phase of three-phase AC voltage to the first group of six racks 606, a second phase of three-phase AC voltage to the second group of six racks 608, and a third phase of three-phase AC voltage to the third group of six racks 610. In other embodiments, the quantity of racks and computing systems can vary.

FIG. 7 shows a control distribution system 700 of the flexible datacenter 500 according to one or more example embodiments. The system 700 includes a grid operator 702, a generation station control system 216, a remote master control system 300, and a flexible datacenter 500. As such, the system 700 represents one example configuration for controlling operations of the flexible datacenter 500, but other configurations may include more or fewer components in other arrangements.

The datacenter control system 504 may independently or cooperatively with one or more of the generation station control system 414, the remote master control system 300, and/or the grid operator 702 modulate power at the flexible datacenter 500. During operations, the power delivery to the flexible datacenter 500 may be dynamically adjusted based on conditions or operational directives. The conditions may correspond to economic conditions (e.g., cost for power, aspects of computational operations to be performed), power-related conditions (e.g., availability of the power, the sources offering power), demand response, and/or weather-related conditions, among others.

The generation station control system 414 may be one or more computing systems configured to control various aspects of a generation station (not independently illustrated, e.g., 216 or 400). As such, the generation station control system 414 may communicate with the remote master

36

control system 300 over a networked connection 706 and with the datacenter control system 704 over a networked or other data connection 708.

As discussed with respect to FIGS. 2 and 3, the remote master control system 300 can be one or more computing systems located offsite, but connected via a network connection 710 to the datacenter control system 504. The remote master control system 300 may provide supervisory controls or override control of the flexible datacenter 500 or a fleet of flexible datacenters (not shown).

The grid operator 702 may be one or more computing systems that are configured to control various aspects of the power grid (not independently illustrated) that receives power from the generation station. The grid operator 702 may communicate with the generation station control system 300 over a networked or other data connection 712.

The datacenter control system 504 may monitor BTM power conditions at the generation station and determine when a datacenter ramp-up condition is met. The BTM power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost to obtain power), low priced power, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs. For example, a datacenter control system may analyze future workload and near term weather conditions at the flexible datacenter.

In some instances, the datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702 to go offline or reduce power. As such, the datacenter control system 504 may enable the power input system 502 to provide power to the power distribution system 506 to power the computing systems 512 or a subset thereof.

The datacenter control system 504 may optionally direct one or more computing systems 512 to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems 512 are configured to perform blockchain hashing operations, the datacenter control system 504 may direct them to perform blockchain hashing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems 512 may be configured to perform high-throughput computing operations and/or high performance computing operations.

The remote master control system 300 may specify to the datacenter control system 504 what sufficient behind-the-meter power availability constitutes, or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 500. In such circumstances, the datacenter control system 504 may provide power to only a subset of computing systems, or operate the plurality of computing systems in a

US 10,608,433 B1

37

lower power mode, that is within the sufficient, but less than full, range of power that is available. In addition, the computing systems 512 may adjust operational frequency, such as performing more or less processes during a given duration. The computing systems 512 may also adjust internal clocks via over-clocking or under-clocking when performing operations.

While the flexible datacenter 500 is online and operational, a datacenter ramp-down condition may be met when there is insufficient or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702. The datacenter control system 504 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by the remote master control system 300 or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently.

An operational directive may be based on current dispatch-ability, forward looking forecasts for when behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the generation station control system 414, the remote master control system 300, or the grid operator 702. For example, the generation station control system 414, the remote master control system 300, or the grid operator 702 may issue an operational directive to flexible datacenter 500 to go offline and power down. When the datacenter ramp-down condition is met, the datacenter control system 504 may disable power delivery to the plurality of computing systems (e.g., 512). The datacenter control system 504 may disable 714 the power input system 502 from providing power (e.g., three-phase nominal AC voltage) to the power distribution system 506 to power down the computing systems 512 while the datacenter control system 504 remains powered and is capable of returning service to operating mode at the flexible datacenter 500 when behind-the-meter power becomes available again.

While the flexible datacenter 500 is online and operational, changed conditions or an operational directive may cause the datacenter control system 504 to modulate power consumption by the flexible datacenter 500. The datacenter control system 504 may determine, or the generation station control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power the flexible datacenter 500. In such situations, the datacenter control system 504 may take steps to reduce or stop power consumption by the flexible datacenter 500 (other than that required to maintain operation of datacenter control system 504).

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically reduce or withdraw power delivery to one or more computing systems 512 to meet the dictate. The datacenter control system 504 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (e.g., 512) to reduce power consumption. The datacenter control system 504 may dynamically reduce the power consumption of one or more computing

38

systems by reducing their operating frequency or forcing them into a lower power mode through a network directive.

Similarly, the flexible datacenter 500 may ramp up power consumption based on various conditions. For instance, the datacenter control system 504 may determine, or the generation control system 414, the remote master control system 300, or the grid operator 702 may communicate, that a change in local conditions may result in greater power generation, availability, or economic feasibility. In such situations, the datacenter control system 504 may take steps to increase power consumption by the flexible datacenter 500.

Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to increase power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically increase power delivery to one or more computing systems 512 (or operations at the computing systems 512) to meet the dictate. For instance, one or more computing systems 512 may transition into a higher power mode, which may involve increasing power consumption and/or operation frequency.

One of ordinary skill in the art will recognize that datacenter control system 504 may be configured to have a number of different configurations, such as a number or type or kind of the computing systems 512 that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available behind-the-meter power. As such, the datacenter control system 504 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

FIG. 8 shows a control distribution system 800 of a fleet of flexible datacenters according to one or more example embodiments. The control distribution system 800 of the flexible datacenter 500 shown and described with respect to FIG. 7 may be extended to a fleet of flexible datacenters as illustrated in FIG. 8. For example, a first generation station (not independently illustrated), such as a wind farm, may include a first plurality of flexible datacenters 802, which may be collocated or distributed across the generation station. A second generation station (not independently illustrated), such as another wind farm or a solar farm, may include a second plurality of flexible datacenters 804, which may be collocated or distributed across the generation station. One of ordinary skill in the art will recognize that the number of flexible datacenters deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more example embodiments.

The remote master control system 300 may provide directive to datacenter control systems of the fleet of flexible datacenters in a similar manner to that shown and described with respect to FIG. 7, with the added flexibility to make high level decisions with respect to fleet that may be counterintuitive to a given station. The remote master control system 300 may make decisions regarding the issuance of operational directives to a given generation station based on, for example, the status of each generation station where flexible datacenters are deployed, the workload distributed across fleet, and the expected computational demand required for one or both of the expected workload and predicted power availability. In addition, the remote master control system 300 may shift workloads from the first plurality of flexible datacenters 802 to the second plurality of flexible datacenters 804 for any reason, including, for

example, a loss of BTM power availability at one generation station and the availability of BTM power at another generation station. As such, the remote master control system 300 may communicate with the generation station control systems 806A, 806B to obtain information that can be used to organize and distribute computational operations to the fleets of flexible datacenters 802, 804.

FIG. 9 shows a queue distribution arrangement for a traditional datacenter 902 and a flexible datacenter 500, according to one or more example embodiments. The arrangement of FIG. 9 includes a flexible datacenter 500, a traditional datacenter 902, a queue system 312, a set of communication links 916, 918, 920A, 920B, and the remote master control system 300. The arrangement of FIG. 9 represents an example configuration scheme that can be used to distribute computing operations using a queue system 312 between the traditional datacenter 902 and one or more flexible datacenters. In other examples, the arrangement of FIG. 9 may include more or fewer components in other potential configurations. For instance, the arrangement of FIG. 9 may not include the queue system 312 or may include routes that bypass the queue system 312.

The arrangement of FIG. 9 may enable computational operations requested to be performed by entities (e.g., companies). As such, the arrangement of FIG. 9 may use the queue system 312 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 500 and the critical traditional datacenter 902. Particularly, the arrangement of FIG. 9 may use the queue system 312 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations among datacenters. As a result, the use of the queue system 312 may reduce time to complete operations and reduce costs.

In some examples, one or more components, such as the datacenter control system 504, the remote master control system 300, the queue system 312, or the control system 936, may be configured to identify situations that may arise where using the flexible datacenter 500 can reduce costs or increase productivity of the system, as compared to using the traditional datacenter 902 for computational operations. For example, a component within the arrangement of FIG. 9 may identify when using behind-the-meter power to power the computing systems 512 within the flexible datacenter 500 is at a lower cost compared to using the computing systems 934 within the traditional datacenter 902 that are powered by grid power. Additionally, a component in the arrangement of FIG. 9 may be configured to determine situations when offloading computational operations from the traditional datacenter 902 indirectly (i.e., via the queue system 312) or directly (i.e., bypassing the queue system 312) to the flexible datacenter 500 can increase the performance allotted to the computational operations requested by an entity (e.g., reduce the time required to complete time-sensitive computational operations).

In some examples, the datacenter control system 504 may monitor activity of the computing systems 512 within the flexible datacenter 500 and use the respective activity levels to determine when to obtain computational operations from the queue system 312. For instance, the datacenter control system 504 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 512 to perform. The various factors may include power availability at the flexible datacenter 500 (e.g., either stored or from a BTM source), availability of the computing systems 512 (e.g., percentage of computing systems avail-

able), type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter 500, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system, among others. The datacenter control system 504 may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems 512 to perform. In such a configuration, the datacenter control system 504 manages the activity of the flexible datacenter 500, including determining when to acquire new sets of computational operations when capacity among the computing systems 512 permit.

In other examples, a component (e.g., the remote master control system 300) within the system may assign or distribute one or more sets of computational operations organized by the queue system 312 to the flexible datacenter 500. For example, the remote master control system 300 may manage the queue system 312, including the distribution of computational operations organized by the queue system 312 to the flexible datacenter 500 and the traditional datacenter 902. The remote master control system 300 may utilize to information described with respect to the Figures above to determine when to assign computational operations to the flexible datacenter 500.

The traditional datacenter 902 may include a power input system 930, a power distribution system 932, a datacenter control system 936, and a set of computing systems 934. The power input system 930 may be configured to receive power from a power grid and distribute the power to the computing systems 934 via the power distribution system 932. The datacenter control system 936 may monitor activity of the computing systems 934 and obtain computational operations to perform from the queue system 312. The datacenter control system 936 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 934 to perform. A component (e.g., the remote master control system 300) within the arrangement of FIG. 9 may assign or distribute one or more sets of computational operations organized by the queue system 312 to the traditional datacenter 902.

The communication link 916 represents one or more links that may serve to connect the flexible datacenter 500, the traditional datacenter 902, and other components within the system (e.g., the remote master control system 300, the queue system 312—connections not shown). In particular, the communication link 916 may enable direct or indirect communication between the flexible datacenter 500 and the traditional datacenter 902. The type of communication link 916 may depend on the locations of the flexible datacenter 500 and the traditional datacenter 902. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

The queue system 312 represents an abstract data type capable of organizing computational operation requests received from entities. As each request for computational operations are received, the queue system 312 may organize the request in some manner for subsequent distribution to a datacenter. Different types of queues can make up the queue system 312 within embodiments. The queue system 312 may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

US 10,608,433 B1

41

In other examples, the queue system 312 may be distributed consisting of multiple queue sub-systems. In the distributed configuration, the queue system 312 may use multiple queue sub-systems to organize different sets of computational operations. Each queue sub-system may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required for performing each set of computational operations. For instance, a first queue sub-system may organize sets of non-intensive computational operations and a second queue sub-system may organize sets of intensive computational operations. In some examples, the queue system 312 may include queue sub-systems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations. In some examples, the queue system 312 may move computational operations between different computing systems or different datacenters in real-time.

Within the arrangement of FIG. 9, the queue system 312 is shown connected to the remote master control system 300 via the communication link 918. In addition, the queue system 312 is also shown connected to the flexible datacenter via the communication 920A and to the traditional datacenter 902 via the communication link 920B. The communication links 918, 920A, 920B may be similar to the communication link 916 and can be various types of communication links within examples.

The queue system 312 may include a computing system configured to organize and maintain queues within the queue system 312. In another example, one or more other components of the system may maintain and support queues within the queue system 312. For instance, the remote master control system 300 may maintain and support the queue system 312. In other examples, multiple components may maintain and support the queue system 312 in a distributed manner, such as a blockchain configuration.

In some embodiments, the remote master control system 300 may serve as an intermediary that facilitates all communication between flexible datacenter 500 and the traditional datacenter 902. Particularly, the traditional datacenter 902 or the flexible datacenter 500 might need to transmit communications to the remote master control system 300 in order to communicate with the other datacenter. As also shown, the remote master control system 300 may connect to the queue system 312 via the communication link 918. Computational operations may be distributed between the queue system 312 and the remote master control system 300 via the communication link 918. The computational operations may be transferred in real-time and mid-performance from one datacenter to another (e.g., from the traditional datacenter 902 to the flexible datacenter 500). In addition, the remote master control system 300 may manage the queue system 312, including providing resources to support queues within the queue system 312.

As a result, the remote master control system 300 may offload some or all of the computational operations assigned to the traditional datacenter 902 to the flexible datacenter 500. This way, the flexible datacenter 500 can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist traditional datacenter 902. The remote master control system 300 may use the queue system 312 to temporarily store and organize the

42

offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter 500) is available to perform them. The flexible datacenter 500 consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to the traditional datacenter 902. The remote master control system 300 may further communicate with the flexible datacenter 500 via communication link 922 and the traditional datacenter 902 via the communication link 924.

FIG. 10A shows method 1000 of dynamic power consumption at a flexible datacenter using behind-the-meter power according to one or more example embodiments. Other example methods may be used to manipulate the power delivery to one or more flexible datacenters.

In step 1010, the datacenter control system, the remote master control system, or another computing system may monitor behind-the-meter power availability. In some embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1020, the datacenter control system or the remote master control system 300 may determine when a datacenter ramp-up condition is met. In some embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the generation station to go offline or reduce power.

In step 1030, the datacenter control system may enable behind-the-meter power delivery to one or more computing systems. In some instances, the remote master control system may directly enable BTM power delivery to computing systems within the flexible system without instructing the datacenter control system.

In step 1040, once ramped-up, the datacenter control system or the remote master control system may direct one or more computing systems to perform predetermined computational operations. In some embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

While operational, the datacenter control system, the remote master control system, or another computing system may receive an operational directive to modulate power consumption. In some embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system or the remote master control system may dynamically reduce power delivery to one or more computing systems or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system or the remote master control system may dynamically adjust power delivery to one or more computing systems to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system may disable power delivery to one or more computing systems.

FIG. 10B shows method 1050 of dynamic power delivery to a flexible datacenter using behind-the-meter power according to one or more embodiments. In step 1060, the datacenter control system or the remote master control system may monitor behind-the-meter power availability. In

US 10,608,433 B1

43

certain embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

In step 1070, the datacenter control system or the remote master control system may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the generation station to go offline or reduce power.

In step 1080, the datacenter control system may disable behind-the-meter power delivery to one or more computing systems. In step 1090, once ramped-down, the datacenter control system remains powered and in communication with the remote master control system so that it may dynamically power the flexible datacenter when conditions change.

One of ordinary skill in the art will recognize that a datacenter control system may dynamically modulate power delivery to one or more computing systems of a flexible datacenter based on behind-the-meter power availability or an operational directive. The flexible datacenter may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter may have a blackout state, where all power consumption, including that of the datacenter control system is halted. However, once the flexible datacenter enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system. Generation station conditions or operational directives may cause flexible datacenter to ramp-up, reduce power consumption, change power factor, or ramp-down.

FIG. 11 illustrates a block diagram of a system for implementing control strategies based on a power option agreement, according to one or more embodiments. The system 1100 represents an example arrangement that includes a control system (e.g., the remote master control system 262), a load (e.g., one or more of the datacenters 1102, 1104, and 1106), and a power entity 1140, which may establish and operate in accordance with a power option agreement. Additional arrangements are possible within examples.

In general, a power option agreement is an agreement between a power entity 1140 associated with the delivery of power to a load (e.g., a grid operator, power generation station, or local control station) and the load (e.g., the datacenters 1102-1106). As part of the power option agreement, the load (e.g., load operator, contracting agent for the load, semi-automated control system associated with the load, and/or automated control system associated with the load) provides the power entity 1140 with the right, but not obligation, to reduce the amount of power delivered (e.g., grid power) to the load up to an agreed amount of power during an agreed upon time interval. In order to provide the power entity 1140 with this option, the load needs to be using at least the amount of power subject to the option (e.g., a minimum power threshold). For instance, the load may agree to use at least 1 MW of grid power at all times during a specified 24-hour time interval to provide the power entity 1140 with the option of being able to reduce the amount of power delivered to the load by any amount up to 1 MW at any point during the specified 24-hour time interval. The load may grant the power entity 1140 with this option in exchange for a monetary consideration (e.g., receive power

44

at a reduced price and/or monetary payment if the option is exercised by the power entity).

The power option agreement may be used by the power entity 1140 to reserve the right to reduce the amount of grid power delivered to the load during a set time frame (e.g., the next 24 hours). For instance, the power entity 1140 may exercise a predefined power option to reduce the amount of grid power delivered to the load during a time when the grid power may be better redirected to other loads coupled to the power grid. As such, the power entity 1140 may exercise power option agreements to balance loads coupled to the power grid. In some embodiments, a power option agreement may also specify other parameters, such as costs associated with different levels of power consumption and/or maximum power thresholds for the load to operate according to.

To illustrate an example, a power option agreement may specify that a load (e.g., the datacenters 1102-1106) is required to use at least 10 MW or more at all times during the next 12 hours. Thus, the minimum power threshold according to the power option agreement is 10 MW and this minimum power threshold extends across the time interval of the next 12 hours. In order to comply with the agreement, the load must subsequently operate using 10 MW or more power at all times during the next 12 hours. This way, the load can accommodate a situation where the power entity 1140 exercises the option. Particularly, exercising the option may trigger the load to reduce the amount of power it consumes by an amount up to 10 MW at any point during the 12 hour interval. By establishing this power option agreement, the power entity 1140 can manipulate the amount of power consumed at the load during the next 12 hours by up to 10 MW if power needs to be redirected to another load or a reduction in power consumption is needed for other reasons.

In the example arrangement of the system 1100 shown in FIG. 11, one or more of the datacenters (e.g., the flexible datacenters 1102, 1104, and the traditional datacenter 1106) may operate as the load that is subject to a power option agreement. As the load that is subject to the power option agreement, the datacenters 1102-1106 may execute control instructions in accordance with power target consumption targets that meet or exceed the minimum power thresholds based on the power option agreement.

As shown in FIG. 11, each datacenter 1102-1106 may include a set of computing systems configured to perform computational operations using power from one or more power sources (e.g., BTM power, grid power, and/or grid power subject to a power option agreement). In particular, the flexible datacenter 1102 includes computing systems 1108 arranged into a first set 1114A, a second set 1114B, and a third set 1114C, the flexible datacenter 1104 includes computing systems 1110 arranged into a first set 1116A, a second set 1116B, and a third set 1118B, and the traditional datacenter 1106 includes computing systems 1112 arranged into a first set 1118A, a second set 1118B, and a third set 1118C. Each set of computing systems may include various types of computing systems that can operate in one or more modes.

The different sets of computing systems as well as the multiple datacenters are included in FIG. 11 for illustration purposes. In particular, the variety of computing systems represent different configurations that a load may take while operating in accordance with a power option agreement, and each configuration (as detailed herein) may include ramping up or down power consumption and transferring and performing computational operations between sets of comput-

US 10,608,433 B1

45

ing systems and/or datacenters. In other examples, the load that is subject to a power option agreement may take on other configurations (e.g., a single datacenter **1102-1106**, and/or a single set of computing systems).

The remote master control system **262** may serve as a control system that can determine performance strategies and provide control instructions to the load (e.g., one or more of the datacenters **1102-1106**). In particular, the remote master control system **262** can monitor conditions in concert with the minimum power thresholds and time intervals (e.g., power option data) set forth in, and/or derived from, one or more power option agreements to determine performance strategies that can enable the load to meet the expectations of the power option agreement(s) while also efficiently using power to accomplish computational operations. In some instances, the remote master control system **262** may also be subject to the power option agreement and may adjust its own power consumption based on the power option agreement (e.g., ramp up or down power consumption based on the defined minimum power thresholds during time intervals).

To establish a power option agreement, the remote master control system **262** (or another computing system) may communicate with the power entity **1140**. For instance, the remote master control system **262** may provide a request (e.g., a signal and/or a bid) to the power entity **1140** and receive the terms of one or more power option agreements, or power option data related to power option agreements (e.g., data such as minimum power thresholds and time intervals, but not all terms contained within a potential power option agreement) in response. In some examples, the remote master control system **262** may evaluate one or more conditions prior to establishing a power option agreement to ensure that the conditions could enable the load (e.g., the datacenters **1102-1106**) to operate in accordance with the power option agreement. For instance, the remote master control system **262** may check the quantity and deadlines associated with computational operations assigned to specific datacenters prior to establishing specific datacenters as a load subject to a power option agreement. In some cases, multiple power option agreements may be established. For example, each datacenter **1102-1106** may be subject to a different power option agreement, which may result in the remote master control system **262** managing the power consumption at each of the datacenters **1102-1106** differently.

Within the system **1100** shown in FIG. **11**, the power entity **1140** may represent any type of power entity associated with the delivery of power to the load that is subject to a power option agreement. For instance, the power entity **1140** may be a local station control system, a grid operator, or a power generation source. As such, the power entity **1140** may establish power option agreements with the loads via communication with the loads and/or the remote master control system **262**. For example, the power entity **1140** may obtain and accept a bid from a load trying to engage in a power option agreement with the power entity **1140**. The power entity **1140** is shown with a power option module **1142**, which may be used to establish power option agreements (e.g., fixed-duration **1144** and/or dynamic **1146**).

Once a power option agreement is established, the remote master control system **262** may obtain power option data from the power entity **1140** (or another source) that specifies the power and time expectations of the power entity **1140**. As shown in FIG. **11**, the power entity **1140** includes a power option module **1142**, which may be used to provide power option data to the remote master control system **262** and/or

46

the datacenters **1102-1106**. In particular, the power option data may specify the minimum power threshold or thresholds associated with one or more time intervals for the load to operate at in accordance with based on the power option agreement. The power option data may also specify other constraints that the load should operate in accordance with.

In some examples, the power option data may also include an indication of a monetary penalty that would be imposed upon the load for failure to operate as agreed upon for the power option agreement. In addition, the power option data may also include an indication of a monetary benefit provided to the load operating at power consumption levels that are in accordance with a power option agreement. For instance, monetary benefits could include reduced prices for power, credits for power, and/or monetary payments. In addition, the power option data may include further constraints upon power use, such as one or more maximum power thresholds and corresponding time intervals for the maximum power thresholds.

In some embodiments, the power entity **1140** may correspond to a qualified scheduling entity (QSE). A QSE may submit bids and offers on behalf of resource entities (REs) or load serving entities (LSEs), such as retail electric providers (REPs). QSEs may submit offers to sell and/or bids to buy power (energy) in the Day-Ahead Market (e.g., the next 24 hours) and the Real-Time Market. As such, the remote master control system **262** or another computing system may communicate with one or more QSEs to engage and control one or more loads in accordance with one or more power option agreements.

In some examples, a power option agreement may take the form of a fixed duration power option agreement **1144**. The fixed duration power option agreement **1144** may specify a set of minimum power thresholds and a set of time intervals in advance for an upcoming fixed duration of time covered by the agreement. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. Examples of such association are provided in FIG. **12**. The fixed duration power option agreement may be established in advanced of the time period covered by the set of time intervals to enable the remote master control system **262** to prepare performance strategies for the load (e.g., the datacenter(s)) associated with the power option agreement. Thus, the remote master control system **262** may evaluate the fixed duration power option and other monitored conditions to determine performance strategies for a set of computing systems (e.g., one or more datacenters) during the different intervals that satisfy the minimum power thresholds.

In other examples, a power option agreement may take the form of a dynamic power option agreement **1146**. For a dynamic power option agreement **1146**, minimum power thresholds may be provided to the remote master control system **262** in real-time (or near real-time). For instance, a dynamic power option agreement may specify that the power entity **1140** may provide adjustments to minimum power thresholds and corresponding time intervals in real-time to the remote master control system **262**. For example, a dynamic power option agreement may provide power option data that specifies a minimum power threshold for immediate adjustments (e.g., for the next hour).

In an embodiment, a dynamic power option agreement **1146** may involve repeat communication between the remote master control system **262** and the power entity **1140**. Particularly, the power entity **1140** may provide signals to the remote master control system **262** that request power consumption adjustments to be initiated at one or more

US 10,608,433 B1

47

datacenters by the remote master control system 262 over short time intervals, such as across minutes or seconds. For example, the power entity 1140 may communicate to the remote master control system 262 to ramp power consumption down to a particular level within the next 5 minutes. As a result, the remote master control system 262 may provide instructions to one or more datacenters to ramp down power consumption using a linear ramp over the next 5 minutes to meet the particular level specified by the power entity 1140. The remote master control system 262 may monitor the linear ramp down of power consumption and increase or decrease the rate that the datacenter(s) ramp down power use based on projections and updates received from the power entity 1140. As a result, although the ramp down of power consumption may initially be performed in a linear manner to meet a power target threshold, the remote master control system 262 may adjust the rate of power consumption decrease based on updates from the power entity 1140. For example, 25 percent of the overall power consumption ramp down may occur during a first period (e.g., 4 minutes 30 seconds) of the 5 minutes and the remaining 75 percent of the overall power consumption ramp down may occur during the remaining period of the 5 minutes (e.g., the final 30 seconds). The example percentages are included for illustration purposes and can vary within examples based on various parameters, such as additional communication (e.g., adjustments) provided by the power entity 1140.

In further examples, a power option agreement may operate similarly to both a fixed-duration 1144 and a dynamic power option agreement 1146. Particularly, power option data specifying minimum power thresholds and corresponding time intervals may be provided in advance for the entire fixed-duration of time (e.g., the next 24 hours). Additional power option data may then be subsequently provided enabling the remote master control system 262 to make one or more adjustments to accommodate any changes specified within the additional power option data. For instance, additional power option data may indicate that a power entity exercised its option to deliver less power to the load. As a result, the remote master control system may instruct the load to adjust power consumption based on the power entity reducing the power threshold minimum via exercising the option.

As indicated above, the remote master control system 262 may monitor conditions in addition to the constraints set forth in power option data received from the power entity 1140. Particularly, the remote master control system 262 may monitor and analyze a set of conditions (including the power option data) to determine strategies for assigning, transferring, and otherwise managing computational operations using the one or more datacenters 1102-1106. The determined strategies may enable efficient operation by the datacenters while also ensuring that the datacenters operate at target power consumption levels that meet or exceed the minimum power thresholds set forth within one or more power option agreements.

Example monitored conditions include, but are not limited to, power availability 1120, power prices 1122, computing systems parameters 1124, cryptocurrency prices 1126, computational operation parameters 1128, and weather conditions 1129. Power availability 1120 may include determining power consumption ranges at a set of computing systems and/or at one or more datacenters. In addition, power availability 1120 may also involve determining the source or sources of power available at a datacenter. For instance, the remote master control system 262 may identify the types of power sources (e.g., BTM, grid

48

power, and/or a battery system) that a datacenter has available. Power prices 1122 may involve an analysis of the different costs associated with powering a set of computing systems. For instance, the remote master control system 262 may determine cost of power from the grid without a power option agreement relative to the cost power from the grid under the power option agreement. In addition, the remote master control system 262 may also compare the cost of grid power relative to the cost of BTM power when available at a datacenter. The power prices 1122 may also involve comparing the cost of using power at different datacenters to determine which datacenter may perform computational operations at a lower cost.

Monitoring computing system parameters 1124 may involve determining parameters related to the computing systems at one or more datacenters. For instance, the remote master control system 262 may monitor various parameters of the computing systems at a datacenter, such as the abilities and availability of various computing systems, the status of the queue used to store computational operations awaiting performance by the computing systems. The remote master control system 262 may determine types and operation modes of the computing systems, including which computing systems could operate in different modes (e.g., a higher power or a lower power mode) and/or at different hash rates and/or frequencies. The remote master control system 262 may also estimate when computing systems may complete current computational operations and/or how many computational operations are assigned to computing systems.

Monitoring cryptocurrency prices 1126 may involve monitoring the current price of one or more cryptocurrencies, the hash rate and/or estimated power consumption associated with mining each cryptocurrency, and other factors associated with the cryptocurrencies. The remote master control system 262 may use data related to monitoring cryptocurrency prices 1126 to determine whether using computing systems to mine a cryptocurrency generates more revenue than the cost of power required for performance of the mining operations.

The remote master control system 262 may monitor parameters related to computational operations (e.g., computational operation parameters 1128). For example, the remote master control system 262 may monitor parameters related to the computational operations requiring performance and currently being performed, such quantity of operations, estimated time to complete, cost to perform each computational operation, deadlines and priorities associated with each computational operation. In addition, the remote master control system 262 may analyze computational operations to determine if a particular type of computing system may perform the computational operation better than other types of computing systems.

Monitoring weather conditions 1129 may include monitoring for any potential power generation disruption due to emergencies or other events, and changes in temperatures or weather conditions at power generators or datacenters that could affect power generation. As such, the operations and environment analysis module (or another component) of the remote master control system 262 may be configured to monitor one or more conditions described above.

The performance strategy determined by the remote master control system 262 based on the monitored conditions and/or power option data can include control instructions for the load (e.g., the datacenters and/or one or more sets of computing systems). For instance, a performance strategy can specify operating parameters, such as operating frequen-

US 10,608,433 B1

49

cies, power consumption targets, operating modes, power on/off and/or standby states, and other operation aspects for computing systems at a datacenter.

The performance strategy can also involve aspects related to the assignment, transfer, and performance of computational operations at the computing systems. For instance, the performance strategy may specify computational operations to be performed at the computing systems, an order for completing computational operations based on priorities associated with the computational operations, and an identification of which computing systems should perform which computational operations. In some instances, priorities may depend on revenue associated with completing each computational operation and deadlines for each computational operation.

The monitored conditions may enable efficient distribution and performance of computational operations among computing systems at one or more datacenters (e.g., datacenters 1102-1106) in ways that can reduce costs and/or time to perform computational operations, take advantage of availability and abilities of computing systems at the datacenters 1102-1106, and/or take advantage in changes in the cost for power at the datacenters 1102-1106. In addition, the monitored conditions may also involve consideration of the power option data to ensure that the computing systems consume enough power to meet minimum power thresholds set forth in one or more power option agreements.

The various monitored conditions described above as well as other potential conditions may change dynamically and with great frequency. Thus, to enable efficient distribution and performance of the computational operations at the datacenters, the remote master control system 262 may be configured to monitor changes in the various conditions to assist with the efficient management and operations of the computing systems at each datacenter. For instance, the remote master control system 262 may engage in wired or wireless communication 1130 with datacenter control systems (e.g., datacenter control system 504) at each datacenter as well as other sources (e.g., the power entity 1140) to monitor for changes in the conditions.

The remote master control system 262 may analyze the different conditions in real-time to modulate operating attributes of computing systems at one or more of the datacenters. By using the monitored conditions, the remote master control system 262 may increase revenue, decrease costs, and/or increase performance of computational operations via various modifications, such as transferring computational operations between datacenters or sets of computing systems within a datacenter and adjusting performance at one or more sets of computing systems (e.g., switching to a low power mode).

In some examples, the traditional datacenter 1106 may be the load subject to a power option agreement. As such, the remote master control system 262 may factor the power option agreement when determining whether to perform computational operations using the computing systems 1112 at the traditional datacenter 1106 and/or transfer computational operations to the computing systems 1108, 1110 at the flexible datacenters 1102, 1104. For instance, the monitored conditions may indicate that the price of grid power is substantially higher than BTM power. As a result, the remote master control system 262 may transfer a subset of computational operations from the traditional datacenter 1106 to the flexible datacenters 1102, 1104. The traditional datacenter 1106 may still have some computational operations to perform to ensure that the traditional datacenter 1106 is

50

using enough power to meet the minimum power threshold or thresholds set forth in the power option agreement.

In some examples, the remote master control system 262 may monitor the grid frequency signal received from the power entity 1140. When the frequency of the grid deviates a threshold amount (e.g., 0.036 Hz above or below 60 Hz), the remote master control system 262 may adjust performance strategies at the load. In some cases, the remote master control system 262 may adjust the power consumption at the load, the number of miners (or computing systems) operating at the load, and/or the frequency or hash rate, among other possible changes. The remote master control system may readjust performance strategies at the load in response to receiving additional power option data from the power entity 1140 (e.g., an indication that the frequency of the grid is back to 60 Hz). In addition, the remote master control system 262 may communicate changes in operations at the load to the power entity 1140. This way, the power entity 1140 may obtain confirmation that the load is adjusting in accordance with a power option agreement.

In some embodiments, a power generation source (e.g., the generation station 400 shown in FIG. 4) may enter into a power option agreement with a grid operator, which may provide the grid operator with the option to reduce the amount of power that the power source generator can deliver to the grid during a defined time interval. For instance, a wind generation farm may enter into the power option agreement with the grid operator. In addition, the remote master control system 262 may also enter into a power option agreement with the power generation source (e.g., the wind farm) to provide a load that can receive excess power from the power generation source when the grid operator exercises the option and lowers the amount of power that the power generation source can deliver to the grid. Thus, rather than reducing the amount of power produced, the power generation source could exercise an option in the agreement with remote master control system 262 and redirect excess power to one or more loads (e.g., a set of computing systems) that could ramp up power consumption in response. In such situations, the remote master control system 262 maybe able to use the excess power from the power generation source (e.g., BTM power) to perform operations at one or more loads at a low cost (or no cost at all). In addition, the power generation source may benefit from the power option agreement by directing excess power to the load instead of temporarily halting power production.

In some examples, a power option agreement may depend on parameters associated balancing grid capacity and demand. For instance, power option agreements may incentivize power consumption ramping during periods of peak grid power use.

FIG. 12 shows a graph representing power option data based on a power option agreement, according to one or more embodiments. The graph 1200 shows power option data arranged according to power 1204 over time 1202. As shown in FIG. 12, time 1202 increases along the X-axis and minimum power thresholds 1204 increase along the Y-axis of the graph 1200. In the example embodiment shown in FIG. 12, the time 1202 increases up to a full day (e.g., 24 hours) in 4 hour increments and the power is shown in MW increasing in intervals of 5 MW. The 24 duration and example minimum power thresholds can differ in other embodiments. Particularly, these values may depend on the terms set forth within the power option agreement.

The graph line 1206 represents sets of minimum power thresholds 1206A, 1206B, 1206C that are specified by

US 10,608,433 B1

51

power option data based on the power option agreement. As shown, the graph line **1206** extends the entire 24 hour duration, which indicates that the set of time intervals associated with minimum power thresholds add up to 24 hours. In other examples, the power option agreement may not include a minimum power threshold during a portion of the duration.

The graph line **1206** of the graph **1200** is further used to illustrate power consumption levels that one or more loads (e.g., a set of computing systems) operating according to the power option agreement may utilize during the 24 hour duration. Particularly, the power quantities above the graph line **1206** represents power levels that the load(s) may consume from the power grid during the 24 hour duration that would satisfy the requirements (i.e., the minimum power thresholds **1206A-1206C**) set forth by the power option agreement. In particular, the power quantities above the graph line **1206** include any power quantity that meets or exceeds the minimum power threshold at that time. By extension, the power quantities positioned below the graph line **1206** represents the amount of power that the load could be directed to reduce power consumption by per the power option agreement.

To further illustrate, an initial minimum power threshold **1206A** is shown associated with the time interval starting at hour 0 and extending to hour 8. In particular, the minimum power threshold **1206A** is set at 5 MW during this time interval. Thus, based on the power option data shown in FIG. **12**, the loads must be able to operate at a target power consumption level that is equal to or greater than the 5 MW minimum power threshold **1206A** at all times during the time interval extending from hour 0 to hour 8, in order to be able to satisfy the power option if it is exercised for that time interval. Similarly, the power entity could reduce the power consumed by loads by any amount up to 5 MW at any point during the time interval from hour 0 to hour 8 in accordance with the power option agreement. For instance, the power entity could exercise its option at any point during this time interval to reduce the power consumed by the loads by 3 MW as a way to load balance the power grid. In response to the power entity exercising its option, the load may then operate using 3 MW less power and/or another strategy determined by a control system factoring additional conditions (e.g., the price of grid power, the revenue that could be generated from mining a cryptocurrency, and/or parameters associated with computational operations awaiting performance).

As further shown in the graph **1200** illustrated in FIG. **12**, the next minimum power threshold **1206B** is associated with the following time interval, which starts at hour 8 and extends until hour 16. During this time interval (hour 8 to hour 16), the load(s) may consume 10 MW or more power since the minimum power threshold **1206B** is now set at 10 MW as shown on the Y-axis of the graph **1200**. In light of the power option data, a control system may determine and provide a performance strategy to the load (e.g., a set of computing systems) that includes a power consumption target that meets or exceeds the minimum power threshold **1206B** (i.e., 10 MW). The performance strategy may depend on the power option data as well as other possible conditions, such as the price of grid power, the availability of computing systems, and/or the type of computing operations, etc. In addition, the power entity could exercise its option to reduce the amount of power consumed by the load by 10 MW or less as represented by the power levels under the minimum threshold **1206B** that extend during the time interval of hour 8 to hour 16.

52

The last minimum power threshold **1206C** is associated with the time interval that starts at hour 16 and extends until hour 24. Similar to the initial minimum power threshold **1206A** associated with the beginning of the graph line **1206**, the last minimum power threshold **1206** is also set at 5 MW. As such, at any point during this interval (hour 16 to hour 24) the loads may consume 5 MW or more to operate in accordance with the power option agreement. As discussed above, by operating at 5 MW or more, the load enables the power consumed from the power grid to be reduced any amount from zero up to 5 MW during this time interval.

When determining the power consumption strategy for a load, a computing system (e.g., the remote master control system **262**) may consider various conditions in addition to the power option data received based on one or more power option agreements. Particularly, the computing system may consider and weigh different conditions in addition to the power option data to determine power consumption targets and/or other control instructions for a load. The conditions may include, but are not limited to, the price of grid power, the price of alternative power sources (e.g., BTM power, stored energy), the revenue associated with mining for one or more cryptocurrencies, parameters related to the computational operations requiring performance (e.g., priorities, deadlines, status of the queue organizing the operations, and/or revenue associated with completing each computational operation), parameters related to the set of computing systems (e.g., types and availabilities of computing systems), and other conditions (e.g., penalties if a minimum power threshold is not met and/or monetary benefits from operating under a power option agreement). By weighing various conditions, the computing system may efficiently manage the set of computing systems, including enabling performance of computational operations cost effectively and/or ensuring that computing systems operate at target power consumption levels that one or more satisfy power option agreements.

In some examples, the computing system may decrease the amount of power that a set of computing systems consumes from one source and while also increasing the amount of power that the set consumes from another source. For instance, the computing system may determine that the price of power grid power is above a threshold price that makes computational operations relatively expensive to perform using grid power. As a result, the computing system may provide control instructions for the computing systems to consume power grid power that matches a minimum power threshold specified by power option data. This may enable the computing systems to satisfy the power option agreement while also avoiding using pricey grid power beyond the minimum amount required per the power option data. In addition, the computing system may instruct some computing systems to switch to a low power mode or temporarily stop until the price of power from the grid decreases. The computing system may instruct one or more computing systems to operate using power from another source (e.g., BTM power and/or stored energy from a battery system) and/or transfer one or more computational operations to another set of computing systems (e.g., a different datacenter).

When the power option agreement is a fixed duration power option agreement, the computing system may receive an indication of all the minimum power thresholds **1206A-1206C** and an indication of the associated time interval altogether and in advance of the duration associated with the power option agreement. By providing all of the minimum power thresholds **1206A-1206C** and the time intervals in

US 10,608,433 B1

53

advance, the computing system may determine a performance strategy for the load that can extend across the entire duration. Particularly, the computing system may factor the minimum power thresholds and associated time intervals as well as other monitored conditions to determine the performance strategy for the total duration. This can enable the computing system to accept and assign computational operations to computing systems in advance while also using a performance strategy that meets the expectations of a power option agreement.

In some examples, the performance strategy determined by the computing system may include control instructions for the set of computing systems to execute if a power option is exercised. For instance, the performance strategy may specify different power consumption targets for the computing systems that depend on whether a power option is exercised during each time interval.

In some instances, the computing system may modify the performance strategy when one or more conditions change enough to warrant a modification. For instance, the computing system may receive an indication of a change in a minimum power threshold (e.g., a decrease in the minimum power threshold) and determine one or more modifications based on the new minimum power threshold and/or other conditions (e.g., a change in the price of power).

In other examples, the power option agreement may be a dynamic power option agreement. Particularly, the load may be subject to a changing minimum power threshold that can vary during a predefined duration associated with the power option agreement. For example, a dynamic power option agreement may specify that the load is subject to a minimum power threshold that may vary from 0 MW up to 5 MW during the next 24 hours and the particular minimum threshold for each hour may depend on power option data received from the power entity during the prior hour. The dynamic power option agreement may further specify the expected response time from the load. For instance, the power option agreement may indicate that an indication of a new minimum power threshold will be provided an hour prior to the start of the minimum power threshold. The computing system, for example, may receive an indication at hour 7 about the increase in the minimum power threshold **1206B** starting at hour 8. The indication may (or may not) specify the total time interval associated with a new minimum power threshold. For instance, the indication received by the computing system may specify that the 10 MW minimum power threshold **1206B** extends from hour 8 until hour 16. In other instances, the power option data may indicate that the computing system should abide by the new minimum power threshold until receiving further power option data indicating a change to another new minimum power threshold.

In some examples, the power option data may arrive at the computing system in an unknown order from the power entity with expectations of swift power consumption adjustments by the load. As a result, the power option agreement may require fast ramping of the load to meet changes. Ramping may involve ramping up or down power consumption as well as ramping operating techniques (e.g., adjusting frequency or operation mode).

In some embodiments, the type of power option power agreement may depend on the delivery and content of power option data provided to the load (or a control system controlling the load). For instance, a computing system may receive minimum power thresholds set across an entire duration associated with a power option agreement in advance when the power option agreement is a fixed-

54

duration power option agreement. In other instances, the computing system may receive power option data dynamically and adjust operations in real-time (or near real-time). For instance, the computing system may receive a series of power option data that each specifies minimum power threshold changes during the duration set forth in the dynamic power option agreement. To illustrate an example, the computing system may receive power option data during hour 1 that specifies the minimum power threshold for hour 2, power option data during hour 2 that specifies the minimum power threshold for hour 3, and so on across the duration of the dynamic power option agreement.

In some examples, the minimum power threshold for a time interval may be zero during the duration of a power option agreement. As such, the load may use any amount of power from the power grid in accordance with the power option agreement, including no power at all during this time interval. When the price for power is high during this time frame, the load may ramp down power usage to zero MW to avoid paying the high price for power while still being in compliance with the power option agreement.

FIG. 13 illustrates a method for implementing control strategies based on a fixed-duration power option agreement, according to one or more embodiments. The method **1300** serves as an example and may include other steps within other embodiments. A control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1300**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At step **1302**, the method **1300** involves monitoring a set of conditions. For instance, a computing system (e.g., a control system) may monitor various conditions that could impact the performance of operations at one or more loads, including the power consumption targets at the loads. The set of monitored conditions may include a variety of information obtained from one or more external sources, such as one or more datacenters, databases, power generation stations, or types of sources.

Some example conditions include, but are not limited to, the price of grid power, the price and availability of alternative power options (e.g. BTM power, and/or stored energy), parameters of the load (e.g., ramping abilities, type of computing systems, operation modes, etc.), parameters of tasks to be performed using the power at the load (e.g., types, deadlines, priorities, and/or revenue associated with computational operations), availability of other computing systems and their associated costs, and/or revenue associated with mining a cryptocurrency. The computing system may monitor one or more of these conditions as well as others.

At step **1304**, the method **1300** involves receiving power option data based, at least in part, on a power option agreement. As discussed above, the computing system (e.g., a remote master control system) may engage in a power option agreement with a power entity. As a result, the computing system may control a load (e.g., a set of computing systems) in accordance with power thresholds and time intervals received from the power entity based on the power option agreement.

In some examples, the power option data may specify a set of minimum power thresholds and a set of time intervals. Each minimum power threshold in the set of minimum power thresholds may be associated with a time interval in the set of time intervals. To illustrate an example, the power option data may specify a first minimum power threshold

US 10,608,433 B1

55

associated with a first time interval and a second minimum power threshold associated with a second time interval, with the second time interval subsequent to the first time interval.

The set of time intervals may add up to the duration represented by the power option agreement. For instance, the total duration of the set of time intervals may correspond to a twenty-four hour period (e.g., the next day). In other examples, the power option agreement may span across a different duration (e.g., 12 hours). In additional embodiments, the power option data may specify other information, such as monetary incentives associated with parameters of the power option agreement and/or one or more maximum power thresholds.

At step **1306**, the method **1300** involves determining a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions. The performance strategy may be determined responsive to receiving the power option data. In addition, the performance strategy may include a power consumption target for the set of computing systems for each time interval in the set of time intervals. In some examples, each power consumption target is equal to or greater than the minimum power threshold associated with each time interval.

As an example, the performance strategy may specify a first power consumption target for the set of computing systems for a first time interval such that the first power consumption target is equal to or greater than a first minimum power threshold associated with the first time interval and a second power consumption target for the set for a second time interval in a similar manner (i.e., the second power consumption target is equal to or greater than a second minimum power threshold).

In some examples, the performance strategy may include an sequence for the set of computing systems to follow when performing computational operations. The sequence, for example, may be based on priorities associated with the computational operations. In addition, the performance strategy may include one or more power consumption targets that are greater than the minimum power thresholds when the price of power from the power grid is below a threshold price during the time intervals associated with the minimum power thresholds.

The performance strategy may also involve transferring, delaying, or adjusting one or more computational operations performed at the set of computing systems. In addition, the performance strategy may involve adjusting operations at the computing systems. For instance, one or more computing systems may switch modes (e.g., operate at a higher frequency or switch to a low power mode).

In addition, the performance strategy may also specify power consumption targets for the set of computing systems to use if the power option is exercised during an interval. This way, the computing systems may continue to perform computational operations (or suspend performance) based on the power option being exercised.

At step **1308**, the method **1300** involves providing instructions to the set of computing systems to perform one or more computational operations based on the performance strategy. For example, the set of computing systems may operate according to the performance strategy to ensure that the minimum power thresholds are met during the defined time intervals based on the power option agreement.

Some examples may further involve receiving subsequent power option data based, at least in part, on the power option agreement. The subsequent power option data may specify to decrease one or more minimum power thresholds of the

56

set of power thresholds. Responsive to receiving the subsequent power option data, the performance strategy for the set of computing systems may be modified based on a combination of at least a portion of the subsequent power option data and one or more conditions of the monitored conditions. The modified performance strategy may include one or more reduced power consumption targets for the set of computing systems. The amount of the reduction in a power consumption target may depend linearly with the amount that the corresponding minimum power threshold was reduced by. For instance, when a minimum power threshold for a time interval is reduced from 10 MW to 5 MW, the power consumption target for that time interval may be reduced from 10 MW to 5 MW. Instructions may be provided to the set of computing systems to perform computational operations based on the modified performance strategy.

FIG. 14 illustrates a method for implementing control strategies based on a dynamic power option agreement, according to one or more embodiments. The method **1400** serves as an example and may include other steps within other embodiments. Similar to the method **1400**, a control system (e.g., the remote master control system **262**) may be configured to perform one or more steps of the method **1400**. As such, the control system may take various forms of a computing system, such as a mobile computing device, a wearable computing device, a network of computing systems, etc.

At block **1402**, the method **1400** involves monitoring a set of conditions. Similar to block **1302** of the method **1300**, a computing system may monitor various conditions to determine instructions for controlling a set of computing systems.

At block **1404**, the method **1400** involves receiving first power option data based, at least in part, on a power option agreement while monitoring the set of conditions. The first power option data may specify a first minimum power threshold associated with a first time interval. For example, the first power option data may specify a minimum power threshold of 10 MW for the next hour, which may start in an hour or less.

The power option agreement may correspond to a dynamic power option agreement in some examples. When managing a load with respect to a dynamic power option agreement, a computing system may receive power option data specifying changes in minimum power thresholds that a load (e.g., the set of computing systems) may be designated to use in the near term (e.g., the next hour). For example, the computing system may receive power option data during each hour of the duration specified by a power option agreement that indicates a minimum power threshold for the next hour.

At block **1406**, the method **1400** involves providing first control instructions for a set of computing systems based on a combination of at least a portion of the first power option data and at least one condition. The first control instructions may be provided responsive to receiving the first power option data.

The first control instructions may include a first power consumption target for the set of computing systems for the first time interval. Particularly, the first power consumption target may be equal to or greater than the first minimum power threshold associated with the first time interval. For example, the first power consumption target may be greater than the first minimum power threshold when a cost of power from the power grid is below a threshold price during the first time interval. In other instances, the first power consumption target may be equal to the first minimum power

US 10,608,433 B1

57

threshold when the cost of power from the power grid is greater than the threshold price.

In some examples, control instructions may specify a sequence for the computing systems to follow when performing computational operations. The sequence may be based on priorities associated with each computational operation.

The first control instructions may be determined based on a combination of the first power option data, the price of power from the power grid, and parameters associated with computational operations to be performed at the set of computing systems.

In some examples, the first control instructions may involve ramping up or down power consumption at the set of computing systems. The power consumption may be ramped up or down based on the first minimum power threshold and one or more other conditions (e.g., the price of power).

At block 1408, the method 1400 involves receiving second power option data based, at least in part, on the power option agreement while monitoring the set of conditions. The computing system may receive the second power option data subsequent to receiving the first power option data. The second power option data may specify a second minimum power threshold associated with a second time interval. For example, the second minimum power threshold may be 7 MW over the duration of the upcoming hour. In other examples, the second minimum power threshold may differ as shown in FIG. 12.

In some instances, the computing system may receive the second power option data during the first time interval such that the second time interval overlaps the first time interval. For instance, the computing system may receive the second power option data to enable real-time adjustments to be made to the power consumed at the set of computing systems.

At block 1410, the method 1400 involves providing second control instructions for the set of computing systems based on a combination of at least a portion of the second power option data and at least one condition. The second control instructions may be provided responsive to receiving the second power option data. The second control instructions may specify a second power consumption target for the set of computing systems for the second time interval. The second power consumption target may be equal to or greater than the second minimum power threshold associated with the second time interval.

In some examples, the computing system may provide a request to a QSE to determine the power option agreement. As such, the computing system may receive power option data (e.g., the first and second power option data) in response to providing the request to the QSE.

The computing system may monitor the price of power from the power grid, and the global mining hash rate and a price for a cryptocurrency (e.g., Bitcoin), among other conditions. The computing system may determine control instructions (e.g., the first and/or second control instructions) based on a combination of power option data, the price of power from the power grid, and the global mining hash rate and the price for the cryptocurrency. For instance, the computing system may cause one or more computing systems (e.g., a subset of computing systems) to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

58

Advantages of one or more embodiments of the present invention may include one or more of the following:

One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive low cost or unutilized power behind-the-meter when it is available.

One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations across multiple datacenters.

One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

One or more embodiments of the present invention allows for management and distribution of computational operations at computing systems across a fleet of datacenters such that the performance of the computational operations take advantages of increased efficiency and decreased costs.

It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

US 10,608,433 B1

59

What is claimed is:

1. A system comprising:
 - a set of computing systems, wherein the set of computing systems is configured to perform computational operations using power from a power grid;
 - a control system configured to:
 - monitor a set of conditions;
 - receive power option data based, at least in part, on a power option agreement, wherein the power option data specify: (i) a set of minimum power thresholds, and (ii) a set of time intervals, wherein each minimum power threshold in the set of minimum power thresholds is associated with a time interval in the set of time intervals;
 - responsive to receiving the power option data, determine a performance strategy for the set of computing systems based on a combination of at least a portion of the power option data and at least one condition in the set of conditions, wherein the performance strategy comprises a power consumption target for the set of computing systems for each time interval in the set of time intervals, wherein each power consumption target is equal to or greater than the minimum power threshold associated with each time interval; and
 - provide instructions to the set of computing systems to perform one or more computational operations based on the performance strategy.
2. The system of claim 1, wherein the control system is configured to monitor the set of conditions comprising:
 - a price of power from the power grid; and
 - a plurality of parameters associated with one or more computational operations to be performed at the set of computing systems.
3. The system of claim 2, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at least the portion option data, the price of power from the power grid, and the plurality of parameters associated with the one or more computational operations.
4. The system of claim 3, wherein the performance strategy further comprises:
 - an order for the set of computing systems to follow when performing the one or more computational operations, wherein the order is based on respective priorities associated with the one or more computational operations.
5. The system of claim 4, wherein the performance strategy further comprises:
 - at least one power consumption target that is greater than a minimum power threshold when the price of power from the power grid is below a threshold price during the time interval associated with the minimum power threshold.
6. The system of claim 1, wherein the control system is further configured to:
 - receive subsequent power option data based, at least in part, on the power option agreement,
 - wherein the subsequent power option data specify to decrease one or more minimum power thresholds of the set of minimum power thresholds.
7. The system of claim 6, wherein the control system is further configured to:
 - responsive to receiving the subsequent power option data,
 - modify the performance strategy for the set of computing systems based on a combination of at least the

60

- portion of the subsequent power option data and at least one condition in the set of conditions,
 - wherein the modified performance strategy comprises one or more reduced power consumption targets for the set of computing systems.
8. The system of claim 7, wherein the control system is further configured to:
 - provide instructions to the set of computing systems to perform the one or more computational operations based on the modified performance strategy.
 9. The system of claim 1, wherein the control system is a remote master control system positioned remotely from the set of computing systems.
 10. The system of claim 1, wherein the control system is a mobile computing device.
 11. The system of claim 1, wherein the control system is configured to receive the power option data while monitoring the set of conditions.
 12. The system of claim 1, wherein the control system is further configured to:
 - provide a request to a qualified scheduling entity (QSE) to determine the power option agreement; and
 - receive power option data in response to providing the request to the QSE.
 13. The system of claim 1, wherein the power option data specify: (i) a first minimum power threshold associated with a first time interval in the set of time intervals, and (ii) a second minimum power threshold associated with a second time interval in the set of time intervals,
 - wherein the second time interval is subsequent to the first time interval.
 14. The system of claim 13, wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems such that the performance strategy comprises:
 - a first power consumption target for the set of computing systems for the first time interval, wherein the first power consumption target is equal to or greater than the first minimum power threshold; and
 - a second power consumption target for the set of computing systems for the second time interval, wherein the second power consumption target is equal to or greater than the second minimum power threshold.
 15. The system of claim 1, wherein a total duration of the set of time intervals corresponds to a twenty-four hour period.
 16. The system of claim 1, wherein the set of conditions monitored by the control system further comprise:
 - a price of power from the power grid; and
 - a global mining hash rate and a price for a cryptocurrency; and
 - wherein the control system is configured to:
 - determine the performance strategy for the set of computing systems based on a combination of at the portion of the power option data, the price of power from the power grid, the global mining hash rate and the price for the cryptocurrency,
 - wherein the performance strategy specifies for at least a subset of the set of computing systems to perform mining operations for the cryptocurrency when the price of power from the power grid is equal to or less than a revenue obtained by performing the mining operations for the cryptocurrency.

US 10,608,433 B1

61

17. A method comprising:
 monitoring, by a computing system, a set of conditions;
 receiving, at the computing system, power option data
 based, at least in part, on a power option agreement,
 wherein the power option data specify: (i) a set of
 minimum power thresholds, and (ii) a set of time
 intervals, wherein each minimum power threshold in
 the set of minimum power thresholds is associated with
 a time interval in the set of time intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.

18. The method of claim 17, wherein determining the
 performance strategy for the set of computing systems
 comprises:
 identifying information about the set of computing sys-
 tems; and
 determining the performance strategy to further comprise
 instructions for at least a subset of the set of computing
 systems to operate at an increased frequency based on
 a combination of at least the portion of the power
 option data and the information about the set of com-
 puting systems.

19. The method of claim 17, further comprising:
 receiving subsequent power option data based, at least in
 part, on the power option agreement, wherein the
 subsequent power option data specify to decrease one
 or more minimum power thresholds of the set of
 minimum power thresholds;

62

responsive to receiving the subsequent power option data,
 modifying the performance strategy for the set of
 computing systems based on a combination of at least
 the portion of the subsequent power option data and at
 least one condition in the set of conditions, wherein the
 modified performance strategy comprises one or more
 reduced power consumption targets for the set of com-
 puting systems; and
 providing instructions to the set of computing systems to
 perform the one or more computational operations
 based on the modified performance strategy.

20. A non-transitory computer readable medium having
 stored therein instructions executable by one or more pro-
 cessors to cause a computing system to perform functions
 comprising:
 monitoring a set of conditions;
 receiving power option data based, at least in part, on a
 power option agreement, wherein the power option
 data specify: (i) a set of minimum power thresholds,
 and (ii) a set of time intervals, wherein each minimum
 power threshold in the set of minimum power thresh-
 olds is associated with a time interval in the set of time
 intervals;
 responsive to receiving the power option data, determin-
 ing a performance strategy for a set of computing
 systems based on a combination of at least a portion of
 the power option data and at least one condition in the
 set of conditions, wherein the performance strategy
 comprises a power consumption target for the set of
 computing systems for each time interval in the set of
 time intervals, wherein each power consumption target
 is equal to or greater than the minimum power thresh-
 old associated with each time interval; and
 providing instructions to the set of computing systems to
 perform one or more computational operations based
 on the performance strategy.

* * * * *

EXHIBIT B

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN
APPLICATION DATA SHEET (37 CFR 1.76)****Title of
Invention**Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration
Power Option Agreement

As the below named inventor, I hereby declare that:

This declaration
is directed to:☒

The attached application, or

☐United States application or PCT international application number _____
filed on _____.

The above-identified application was made or authorized to be made by me.

I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.

I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001
by fine or imprisonment of not more than five (5) years, or both.**WARNING:**

Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.

LEGAL NAME OF INVENTOR

Inventor: Michael T. McNamara

Date (Optional): Dec 2, 2019

Signature: 

Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

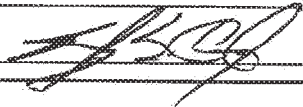
If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

PTO/AIA/01 (06-12)

Approved for use through 01/31/2014. OMB 0651-0032
U.S. Patent and Trademark Office; U.S. DEPARTMENT OF COMMERCE

Under the Paperwork Reduction Act of 1995, no persons are required to respond to a collection of information unless it displays a valid OMB control number.

**DECLARATION (37 CFR 1.63) FOR UTILITY OR DESIGN APPLICATION USING AN
APPLICATION DATA SHEET (37 CFR 1.76)**

Title of Invention	Methods and Systems for Adjusting Power Consumption based on a Fixed-Duration Power Option Agreement	
<p>As the below named inventor, I hereby declare that:</p> <p>This declaration is directed to: <input checked="" type="checkbox"/> The attached application, or <input type="checkbox"/> United States application or PCT international application number _____ filed on _____</p> <p>The above-identified application was made or authorized to be made by me.</p> <p>I believe that I am the original inventor or an original joint inventor of a claimed invention in the application.</p> <p>I hereby acknowledge that any willful false statement made in this declaration is punishable under 18 U.S.C. 1001 by fine or imprisonment of not more than five (5) years, or both.</p> <p style="text-align: center;">WARNING:</p> <p>Petitioner/applicant is cautioned to avoid submitting personal information in documents filed in a patent application that may contribute to identity theft. Personal information such as social security numbers, bank account numbers, or credit card numbers (other than a check or credit card authorization form PTO-2038 submitted for payment purposes) is never required by the USPTO to support a petition or an application. If this type of personal information is included in documents submitted to the USPTO, petitioners/applicants should consider redacting such personal information from the documents before submitting them to the USPTO. Petitioner/applicant is advised that the record of a patent application is available to the public after publication of the application (unless a non-publication request in compliance with 37 CFR 1.213(a) is made in the application) or issuance of a patent. Furthermore, the record from an abandoned application may also be available to the public if the application is referenced in a published application or an issued patent (see 37 CFR 1.14). Checks and credit card authorization forms PTO-2038 submitted for payment purposes are not retained in the application file and therefore are not publicly available.</p>		
<p>LEGAL NAME OF INVENTOR</p> <p>Inventor: <u>Raymond E. Cline Jr.</u> Date (Optional): <u>12/3/2019</u></p> <p>Signature: </p>		
<p>Note: An application data sheet (PTO/SB/14 or equivalent), including naming the entire inventive entity, must accompany this form or must have been previously filed. Use an additional PTO/AIA/01 form for each additional inventor.</p>		

This collection of information is required by 35 U.S.C. 115 and 37 CFR 1.63. The information is required to obtain or retain a benefit by the public which is to file (and by the USPTO to process) an application. Confidentiality is governed by 35 U.S.C. 122 and 37 CFR 1.11 and 1.14. This collection is estimated to take 1 minute to complete, including gathering, preparing, and submitting the completed application form to the USPTO. Time will vary depending upon the individual case. Any comments on the amount of time you require to complete this form and/or suggestions for reducing this burden, should be sent to the Chief Information Officer, U.S. Patent and Trademark Office, U.S. Department of Commerce, P.O. Box 1450, Alexandria, VA 22313-1450. DO NOT SEND FEES OR COMPLETED FORMS TO THIS ADDRESS. **SEND TO: Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.**

If you need assistance in completing the form, call 1-800-PTO-9199 and select option 2.

IN THE UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF DELAWARE

BEARBOX LLC and AUSTIN STORMS,

Plaintiffs,

V.

LANCIUM LLC, MICHAEL T.
MCNAMARA, and RAYMOND E. CLINE,
JR.

Defendants.

C.A. No. 21-534-MN-CJB

JURY TRIAL DEMANDED

SECOND AMENDED COMPLAINT

Plaintiffs BearBox LLC (“BearBox”) and Austin Storms (collectively, “Plaintiffs” or “[BearBox](#)”) bring this action against Lancium LLC (“Lancium”), Michael T. McNamara, and Raymond E. Cline, Jr. (collectively “Defendants”) to correct the inventorship of U.S. Patent No. 10,608,433 (the “433 Patent”) and to recover damages, injunctive relief, declaratory relief, and other remedies for Defendants’ wrongful actions to obtain, [convert](#), misuse, disclose, and claim as their own Plaintiffs’ proprietary cryptocurrency mining technology. Plaintiffs further allege as follows:

INTRODUCTION

1. This case is about the Defendants' theft ~~of and unauthorized use of system designs, data, know-how, and physical documents describing those designs, data, and know-how, as well as~~ inventions that rightfully belong to Plaintiffs.
2. Plaintiffs developed proprietary technology relating to cryptocurrency mining systems ~~(the "BearBox Technology").~~ By way of background, ~~the BearBox Technology~~ BearBox's technology generally relates to an energy-efficient cryptocurrency mining

{01772484;v1 }

Style Definition: Heading 1: Indent: Left: 0", Tab stops: Not at 1.38"

Style Definition: Comment Text

system and related methods that reduce the inefficiency and environmental impact of energy-
~~expensive~~intensive mining operations by better utilizing available energy resources to increase
~~the~~ stability of the energy grid, minimize a mining operation's impact on peak-demand, and also
 alleviate energy ~~over supply~~oversupply and undersupply conditions. ~~The BearBox~~
~~Technology~~BearBox's technology can be used to mine cryptocurrency, such as Bitcoin.

3. The Defendants deceptively induced the Plaintiffs to disclose ~~the BearBox~~
~~Technology~~BearBox's technology to them under the guise of a possible business deal or
"collaboration," a word used by McNamara in encouraging Plaintiffs' disclosures, between
 Defendants and Plaintiffs to jointly commercialize the BearBox Technology. ~~Before disclosing~~
~~the BearBox Technology to Defendants, Plaintiffs obtained assurances of confidentiality from~~
~~Defendants.~~

4. The Defendants stole ~~the BearBox Technology~~BearBox's technology, including
system designs, documents, data, and know-how, from Plaintiffs by converting ~~and~~
~~misappropriating it and claiming using it to modify Defendants' software, e.g. Smart Response™,~~
to operate as their own. ~~BearBox's technology did, and then selling, licensing, seeking~~
investments in support of, and otherwise monetizing that modified Smart Response™ software
for great profit.

5. Defendants ~~filed~~went even further, by filing a U.S. patent application, which later
matured into the '433 Patent, that wrongfully disclosed ~~the BearBox Technology~~portions of
BearBox's technology to the U.S. Patent and Trademark Office and ultimately to the public; as
technology that was invented by Defendants. The claimed subject matter of the '433 Patent falls
 fully within the scope of ~~the BearBox Technology.~~ And by BearBox's technology, though is not a
complete representation of the BearBox technology. By obtaining the '433 Patent with claims

directed to ~~the BearBox Technology~~ portions of BearBox's technology, the Defendants have wrongfully obtained a patent ~~covering the BearBox Technology and wrongfully claimed the BearBox Technology as that~~ prevents Plaintiffs from building their own system, or otherwise monetizing their ~~own~~ system designs, data, documents, and know-how.

6. ~~Plaintiffs~~ Plaintiffs bring this action to recover damages caused by Defendants' theft and unauthorized use, and subsequent exploitation of Plaintiffs system design, data, documents, and know-how.

5-7. ~~Plaintiffs also~~ bring this action to correct the named inventors on the '433 Patent.

The inventions claimed in the '433 Patent are inventions conceived by Storms, founder and president of BearBox.

PARTIES

6-8. ~~Plaintiff~~ BearBox LLC (~~"BearBox"~~) is a limited liability company organized and existing under the laws of Louisiana with its principal place of business at 4422 Highway 22, Mandeville, Louisiana 70471.

7-9. Plaintiff Austin Storms is an individual residing in Mandeville, Louisiana.

8-10. On information and belief, Defendant Lancium is a Delaware limited liability company with its principal place of business at 6006 Thomas Rd, Houston, Texas 77041. On information and belief, Lancium has a registered agent capable of accepting service in this district, Harvard Business Services, Inc. with a place of business at 16192 Coastal Highway, Lewes, DE 19958.

9-11. On information and belief, Defendant Michael T. McNamara is the Chief Executive Officer and a founder of Lancium and resides in Newport Beach, California. Defendant McNamara is named as a purported inventor on the face of the '433 Patent.

~~10-12.~~ On information and belief, Defendant Raymond E. Cline, Jr. is the Chief Computing Officer of Lancium and resides in Houston, Texas. Defendant Cline is named as a purported inventor on the face of the '433 Patent.

JURISDICTION

~~11-13.~~ This is an action seeking correction of the named inventors of a United States patent under 35 U.S.C. § 256. As such, this action arises under the laws of the United States.

~~12-14.~~ This Court has exclusive subject matter jurisdiction under 28 U.S.C. §§ 1331 and 1338(a) because the matter arises under an Act of Congress relating to patents, specifically 35 U.S.C. § 256.

;

~~13-15.~~ The Court has supplemental jurisdiction under 28 U.S.C. § 1367 over all asserted claims under state law because those claims are so related to the claims in this action that arise under federal law that they form part of the same case or controversy.

~~14-16.~~ The Court also has jurisdiction pursuant to 28 U.S.C. § 1332, as complete diversity of citizenship exists among the parties, and the amount in controversy exceeds \$75,000. Plaintiff BearBox is a citizen of the State of Louisiana because it is organized under the laws of the State of Louisiana and has its principal place of business in the State of Louisiana. Plaintiff Storms is a citizen of the State of Louisiana because he resides in the State of Louisiana. In contrast, none of the Defendants are citizens of the State of Louisiana. Defendant Lancium is a citizen of the States of Delaware and Texas because it is organized under the laws of the State of Delaware and has its principal place of business in the State of Texas. Defendant McNamara is a citizen of the State of California because he resides in the State of California. Defendant Cline is a citizen of the State of Texas because he resides in the State of Texas. Therefore, because the

Plaintiffs are both citizens of the State of Louisiana (and no other states) for purposes of diversity jurisdiction, and none of the Defendants are citizens of the State of Louisiana, complete diversity exists among the parties.

~~15.~~17. This Court has general personal jurisdiction over Lancium because it is organized under the laws of the State of Delaware and because it maintains an ongoing presence in this District at least through its registered agent.

~~16.~~18. This Court has specific personal jurisdiction over each of Defendants McNamara and Cline at least under Title 6 of the Delaware Code, § 18-109(a).

~~17.~~19. On information and belief, Defendant McNamara is the Chief Executive Officer of Lancium. On information and belief, as the Chief Executive Offer, McNamara participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

~~18.~~20. McNamara is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from the interests of Lancium and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claims against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant McNamara together. Plaintiffs' claims against Defendant McNamara arise out of his exercise of his powers as Chief Executive Officer of Lancium.

~~19.~~21. On information and belief, Defendant Cline is the Chief Computing Officer of Lancium. On information and belief, as the Chief Computing Officer, Cline participates materially in the management of Lancium, has control and/or decision-making authority over Lancium, and is a key individual who takes actions on behalf of Lancium.

~~20-22.~~ Cline is a necessary or proper party to this action because he has a legal interest in the dispute that is separate from Lancium's interest and because Plaintiffs' claims against him arise out of the same facts and occurrences as the claim against Lancium. Accordingly, it serves judicial economy to consider the claims against Lancium and Defendant Cline together. Plaintiffs' claims against Defendant Cline arise out of his exercise of his powers as Chief Computing Officer of Lancium.

~~21-23.~~ The actions of Defendants McNamara and Cline establish sufficient minimum contacts with Delaware under Delaware law and the United States Constitution to give this Court personal jurisdiction over each of them.

~~22-24.~~ As described below, each Defendant has committed acts giving rise to this action.

VENUE

~~23-25.~~ Venue is proper in this District under 28 U.S.C. § 1391(b)(3) because there is no district in which an action may otherwise be brought as provided in § 1391(b) and Defendant Lancium is subject to the Court's personal jurisdiction with respect to this action.

PLAINTIFFS' PROPRIETARY CRYPTOCURRENCY MINING TECHNOLOGY

~~24-26.~~ As of 2018, the amount of energy required to process computer algorithms to mine cryptocurrencies like Bitcoin was three times greater than the energy required to physically mine gold. Conventional mining of "copper, gold, platinum, and rare earth oxides are 4, 5, 7, and 9 megajoules to generate one U.S. dollars," while "it costs an average of 17 megajoules to mine \$1 worth of bitcoin."¹ The large amount of energy required to mine cryptocurrencies can make

¹ <https://www.marketwatch.com/story/mining-bitcoin-is-3-times-more-expensive-than-mining-gold-research-paper-finds-2018-11-06>

such mining financially prohibitive, and even when financially lucrative, the ~~large~~ energy ~~requirements make~~required for cryptocurrency mining can be harmful to the ~~global~~ environment, with studies showing potential carbon dioxide emissions from cryptocurrency mining “single-handedly rais[ing] global temperatures by 2 degrees by 2023.” *Id.*

~~25-27.~~ At the same time, some forms of electrical power generation, such as wind, solar, or other types of variable renewable energy sources, are ~~terribly~~ inefficient. ~~When producers of electrical power are unable to quickly adjust their operations and suffer from a timing imbalance between peak demand and energy production. Additionally, these non-dispatchable renewable energy sources cannot be controlled by operators~~ in response to dynamically changing grid conditions, ~~these producers~~. As a result, these renewable energy sources and their operators frequently sell power at low or even negative prices, sometimes even being curtailed or shut down entirely, until demand and market prices increase increases and grid conditions change.

~~26-28.~~ Because cryptocurrency mining using proof-of-work is a computationally demanding process, it requires a significant amount of energy to operate. As a result, industrial-scale cryptocurrency mining places a large energy burden on the power grid, driving demand and costs as well as increasing the likelihood of transmission or other grid component failure. ~~if not managed correctly.~~

~~27-29.~~ In late 2018 and early 2019, Austin Storms sought to address these problems by developing energy-efficient cryptocurrency mining systems and methods that reduce the environmental impact of energy-intensive mining operations. Storms conceived of a system that better uses available energy resources to increase the stability of the energy grid, minimize a mining operation’s impact on peak-demand, and help alleviate energy ~~over supply~~oversupply

and undersupply conditions, all while decreasing the overall energy costs of the mining operation and increasing its profitability.

~~28-30.~~ Austin Storms conceived of and developed ~~the BearBox Technology~~ BearBox's technology. Storms is the president and founder of BearBox. ~~The BearBox Technology~~ BearBox's technology includes hardware and software components. Structurally, ~~the BearBox Technology~~ BearBox's technology includes a housing for a plurality of miners (such as ASICs, graphics cards, or the like) under the direction of a smart controller(s).

~~29. — The smart controller monitors various external factors, such as current and expected energy demand and pricing information, current and expected cryptocurrency pricing, and the like. Based on these external factors, the system may determine whether conditions are appropriate to mine cryptocurrency and, if so, subsequently mines the cryptocurrency. Optionally, the system also includes other components for cooling, air filtration, and related features.~~

~~30-31.~~ In ~~the BearBox Technology~~ BearBox's technology, a controller (such as a power distribution unit, network interface, computer, and/or the like) monitors various external factors, such as current and expected energy demand/pricing information, current and expected cryptocurrency pricing, current and expected network hashrate and difficulty, and the like. Based on these external factors, the controller(s) determines appropriate times to mine cryptocurrency in accordance with a desired performance strategy (for example, profitability thresholds). At the appropriate times, the controller initiates mining, for example, by powering on the miners number of miners based on various conditions internal and external conditions.

**DEFENDANTS WRONGFULLY CLAIM THE
BEARBOX'S TECHNOLOGY AS THEIR OWN**

~~31-32.~~ In May 2019, Storms attended the Fidelity FCAT Mining Summit in Boston, Massachusetts on behalf of BearBox to promote ~~the BearBox Technology~~ BearBox's technology and seek potential customers for his revolutionary system.

~~32-33.~~ While at the conference, Storms met Defendant McNamara. Defendant McNamara showed immediate interest in ~~the BearBox Technology~~ BearBox's technology. Under the rouse of a potential business relationship, McNamara deceptively pumped Storms for details about ~~the BearBox Technology~~ BearBox's technology over the course of several exchanges, which included conversations, emails, and text messages about ~~the BearBox Technology~~ BearBox's technology. Storms took McNamara to dinner where McNamara continued to pump Storms for details about ~~the BearBox Technology~~. ~~At all times before and during Storms's disclosure of this information, Storms told McNamara that the BearBox Technology was confidential, and Storms relied on McNamara's good faith assurances that he would keep confidential the information he received from Storms about the BearBox Technology~~ BearBox's technology.

~~33-34.~~ Following the conference, McNamara continued to press Storms for additional details about ~~the BearBox Technology~~ BearBox's technology via text messaging and email. ~~Again relying on Defendant McNamara's assurances of confidentiality,~~ Storms described BearBox's technology and provided annotated system diagrams, component specifications, and modeled data sets to mimic real-world Bitcoin variables such as Bitcoin price and network hashrate, energy prices, time intervals, and power thresholds, and computed profitability figures. Storms communications to McNamara about BearBox's technology was for the sole purpose of

potential partnership. Storms included express confidentiality notices in his email communications with Defendant McNamara.

34-35. After Storms disclosed ~~the BearBox Technology~~ BearBox's technology to McNamara, McNamara abruptly ended all communications with Storms.

35-36. Storms last communicated with McNamara on May 9, 2019 via e-mail, and after sending that message, Storms did not hear from McNamara again.

36-37. At that time, Storms understood that McNamara was not interested in investing in ~~the BearBox Technology.~~ Heor collaborating with BearBox's technology. But Storms had no reason to suspect that ~~McNamara-Defendants were being deceptive and~~ would steal ~~the BearBox Technology and claim~~ BearBox's technology, including the system designs, diagrams, data, documents, and know-how, and use it as his to build Defendants' own system that would function as BearBox's technology did.

38. In the days, weeks, and months following Plaintiffs' disclosures to McNamara, Defendants began working with partners, outside consultants, and vendors to help Defendants bring BearBox's technology to market, labeled as Defendants' own system. Defendants efforts included working internally to modify their Smart Response™ software to work as BearBox's technology did.

39. Defendants soon became confident they could exploit BearBox's technology for profit by modifying their Smart Response™ software, and then using, selling, licensing, and otherwise monetizing that modified Smart Response™ software, but Defendants also realized that others would eventually learn of Defendants' new-and-improved Smart Response™ software, now with the benefits of BearBox's technology, and may copy it in order to compete with Defendants.

40. Defendants believed that the profits they would realize from exploiting their modified Smart Response™ software would be significantly reduced if outside competitors were allowed to copy Defendants' unauthorized use of BearBox's technology and compete directly with Defendants.

41. Because of these realizations, Defendants hatched a plan to patent some of the profitable attributes of BearBox's technology. If such a patent was obtained, Defendants believed that such a patent would effectively grant Defendants a monopoly to exploit significant portions of BearBox's technology for profits that would far exceed any profits should Defendants be left to compete with others, including BearBox.

37-42. On information and belief, Defendants filed U.S. provisional patent application No. 62/927,119 on October 28, 2019, naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

38-43. In addition to falsely claiming to be the inventors of the inventions disclosed in the application, Defendants wrongfully disclosed, without authorization, ~~the confidential BearBox Technology~~ portions of BearBox's technology to the United States Patent and Trademark Office.

39-44. Likewise, on December 4, 2019, Defendants filed U.S. Patent Application Serial No. 16/702,931, once again naming Defendants McNamara and Cline as the purported sole joint inventors of the inventions disclosed in the application.

40-45. The '433 Patent issued on March 31, 2020 naming Defendants McNamara and Cline as the sole purported inventors on the face of the patent. A true and correct copy of the '433 Patent is attached hereto as Exhibit A.

41-46. The inventions claimed in the '433 patent fall within the scope of ~~the BearBox Technology~~BearBox's technology, yet Defendants falsely identified themselves as the inventors of the claimed inventions, when, in fact, Storms is the sole inventor of the claimed inventions. Not all aspects of BearBox's technology that was stolen and used by Defendants was described and claimed in the '433 Patent. For example, Defendants theft and unauthorized use included confidential aspects of BearBox technology's system design, diagrams, data, and know-how to modify Defendants' Smart Response software to develop and exploit, among other things, methods for energy value arbitrage, including cryptocurrency mining systems and arbitraging related energy values.

42-47. On information and belief, McNamara and Cline assigned their purported rights in the '433 patent to Lancium. On information and belief, at all times, Lancium was aware that McNamara and Cline, both officers of Lancium, were not the rightful inventors of ~~the BearBox Technology~~BearBox's technology disclosed in the patent and the inventions claimed in the patent.

43-48. Defendants McNamara and Cline each submitted signed declarations falsely swearing that they were "an original joint inventor" of the claimed subject matter-. A true and correct copy of Defendant McNamara's and Defendant Cline's declarations are attached as Exhibit B.

49. On June 19, 2020, Defendants announced that its Smart Response™ software allowed it to qualify as the "first successful load-only Controllable Load Resource (CLR) designation," which Defendants described as a "breakthrough achievement," that allows for dispatching "excess energy into the grid during times when energy is most expensive," adding

that “with this revolutionary advancement, data centers will also be able to earn additional revenue providing ancillary services to ERCOT.”

50. Since then, Defendants have used BearBox’s technology, including the stolen system designs, diagrams, data, and know-how, and its subsequently-modified Smart Response™ software to function as it did, to allow Defendants to profit significantly from use, license, sale, and investments related to its modified Smart Response™ software. Defendants are not, and have never, been entitled to these profits.

51. In addition to illegally profiting from its theft and subsequent use of BearBox’s technology, Defendants made use of its ill-gotten ’433 patent.

44-52. On August 14, 2020, Lancium filed a lawsuit in the U.S. District Court for the Western District of Texas against Layer1 Technologies, Inc. (“Layer1”) asserting that Layer1 infringes the ’433 patent, and stating that “Lancium’s Controllable Load Resource Technology attracted the interest of many companies, including, upon information and belief, Layer1.” That case is captioned *Lancium LLC v. Layer1 Technologies, Inc.*, Case No. 6:20-cv-739 (W.D. Texas) (the “Layer1 Lawsuit”).

45-53. As part of the Layer1 Lawsuit, Defendants falsely asserted that McNamara and Cline are the sole inventors of the inventions claimed in the ’433 patent.

46-54. Plaintiffs became aware of Defendants’ wrongful use of ~~the BearBox Technology~~BearBox’s technology on or about August 17, 2020, when they learned about the Layer1 Lawsuit through a press release dated August 14, 2020, posted by Lancium on PRNewswire. That press release is available at the following URL:
<https://www.prnewswire.com/news-releases/controllable-load-resource-clr-market-leader-lancium-files-patent-infringement-lawsuit-against-layer1-301112687.html>.

~~47-55.~~ Before seeing the August 14, 2020 press release, Plaintiffs were unaware of Defendants' wrongful use of ~~the BearBox Technology~~ BearBox's technology and was unaware of the '433 patent.

~~48-56.~~ On March 5, 2021, Lancium and Layer1 entered a Stipulation to Dismiss with Prejudice in the Layer1 Lawsuit. According to the stipulation, the parties had entered a Settlement Agreement to resolve the Layer1 Lawsuit.

~~49-57.~~ According to a press release issued by Lancium on March 8, 2021, Lancium and Layer 1 "have entered into a mutually beneficial partnership. Layer1 has licensed Lancium's intellectual property and Lancium will provide Smart Response™ software and services to Layer1." The press release is available at the following URL: <https://www.prnewswire.com/news-releases/lancium-and-layer1-settle-patent-infringement-suit-301242602.html>.

~~50. — On information and belief, as part of the Settlement Agreement between Lancium and Layer1 to settle the Layer1 Lawsuit, Lancium received and continues to receive valuable consideration from Layer1, all of which rightly belongs to Plaintiffs, the rightful owners of the inventions claimed in the '433 Patent.~~

**COUNT I
CORRECTION OF INVENTORSHIP FOR THE '433 PATENT:
AUSTIN STORMS AS SOLE INVENTOR**

~~51-58.~~ Plaintiffs incorporate the above paragraphs by reference.

~~52-59.~~ Storms is the sole inventor of the subject matter claimed in the '433 Patent.

~~53-60.~~ Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the '433 patent and the currently listed inventors on the '433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms or BearBox.

~~54-61.~~ Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the '433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT II
IN THE ALTERNATIVE, CORRECTION OF INVENTORSHIP FOR THE '433
PATENT: AUSTIN STORMS AS JOINT INVENTOR WITH THE CURRENTLY
NAMED INVENTORS

~~55-62.~~ Plaintiffs incorporates the above paragraphs by reference.

~~56-63.~~ In the alternative, Storms is a joint inventor of the subject matter claimed in the '433 Patent and should be added to the individuals currently named as inventors on the '433 Patent.

~~57-64.~~ Through omission, inadvertence, and/or error, Storms was not listed as an inventor on the '433 patent and the currently listed inventors on the '433 patent were improperly listed. The omission, inadvertence, and/or error occurred without any deceptive intent on the part of Storms.

~~58-65.~~ Unless Defendants Lancium, McNamara, and Cline are enjoined from asserting that McNamara and Cline are the sole inventors of the '433 Patent in violation of U.S. federal patent laws, Plaintiffs will suffer irreparable injury. Plaintiffs have no adequate remedy at law.

COUNT III
TRADE SECRET MISAPPROPRIATION UNDER
FEDERAL DEFEND TRADE SECRETS ACT, 18 U.S.C. § 1836

~~66.~~ Plaintiffs incorporate the above paragraphs by reference.

~~67.~~ ~~COUNT III~~ At all relevant times, Plaintiffs and Defendants were engaged in interstate commerce. For example, Plaintiffs and Defendants are engaged in trade and business.

developing and testing, among other things, cryptocurrency mining systems and related methods in interstate commerce throughout the United States.

68. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about BearBox’s technology relating to certain features, including methods for energy value arbitrage, including cryptocurrency mining systems and arbitraging related energy values. Such information constitutes trade secrets of substantial economic value, at least because methods for energy value arbitrage may be used for substantial profit. This confidential technical information was not known to the public or to other persons who can obtain economic value from its disclosure or use. This information constituted a trade secret under the Defend Trade Secrets Act, 18 U.S.C. §§ 1836.

69. Defendants knew the information they received from Austin Storms about BearBox’s technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship, and was in no way the information received to be used by Defendants’ for their own profit. McNamara himself referred to the exchanges with Plaintiffs’ “lots of stuff to collaborate on.” Defendants further knew the technical information they received from Storms about BearBox’s technology was confidential based upon Defendants’ own actions taken to protect that information once they possessed it, and misappropriated it by using in modification of their Smart Response™ software. Defendants’ restricted access to the information, including as it was represented in the Smart Response™ software, required confidentiality agreements for third party customers, potential customers, and licensees, and other internal measures were taken.

70. Plaintiffs included a confidentiality designation on communications to which the confidential trade secret information was attached. Plaintiffs took additional reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of individuals that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its computers. The confidential information was, at all other times, secured on a password protected computer. The confidential information was not externally accessible through the internet. When that password protected computer was not attended, it was secured in a locked building at Storms residence, which included other security measures such as a concierge and elevator keys.

71. For the very limited number of additional occasions the confidential, trade secret information was shared, it was shared with a person or entity subject to contractual obligations regarding confidentiality and/or prohibitions on unauthorized use. For the limited number of occasions BearBox shared the confidential, trade secret information with potential customers, it was done only with BearBox management approval.

72. Defendants misappropriated Plaintiffs' trade secrets when they used BearBox technology, without Plaintiffs' authorization, in at least its Smart Response™ software. Defendants have disclosed and/or used the confidential information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, and using it as Defendants' own, and without permission was improper. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

73. As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages or other financial harm in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, financial loss for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability.

74. The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D).

COUNT IV
TRADE SECRET MISAPPROPRIATION UNDER
LOUISIANA CIVIL PRACTICE AND REMEDIES CODE, TITLE 51, SECTION 13-A

75. Plaintiffs incorporate the above paragraphs by reference.

76. At all relevant times, Plaintiffs and Defendants were engaged in trade or commerce within the meaning of La. R.S. 51:1431-39 (2011), the Louisiana Uniform Trade Secrets Act. For example, Plaintiffs and Defendants are engaged in trade and business, developing and testing, among other things, cryptocurrency mining systems and related methods.

77. In the course of the interactions between Austin Storms and Defendants, Defendants obtained confidential information about BearBox's technology relating to certain features, such as methods for energy value arbitrage, including cryptocurrency mining systems and arbitrating related energy values. Such information constitutes trade secrets of substantial economic value, at least because methods for energy value arbitrage, including cryptocurrency mining systems and arbitrating related energy values may be used for substantial profit. This confidential technical information was not known to the public or to other persons who can

obtain economic value from its disclosure or use. This information constituted a trade secret under the Louisiana Uniform Trade Secret Act.

78. Defendants knew the technical information they received from Austin Storms about BearBox's technology was confidential, as Storms specifically informed Defendant McNamara that the information given to him was to be held in confidence, and was not to be used for any other purpose beyond those necessary to carry out an evaluation in advance of a potential business relationship, and was in no way the information received to be used by Defendants' for their own profit. Defendants further knew the information they received from Storms about BearBox's technology was confidential based upon Defendants' own actions taken to protect that information once they possessed it. Defendants' restricted access to the information, required confidentiality agreements for third party customers, potential customers, and licensees, and other internal measures were taken. Defendants further took these actions to protect Plaintiffs' confidential information after Defendants misappropriated it by including it for use in Defendants' Smart Response™ software.

79. Plaintiffs included a confidentiality designation on communications to which the confidential trade secret information was attached. Plaintiffs took additional reasonable steps to preserve secrecy via these explicit guidelines regarding confidentiality and by limiting the number of collaborators that Plaintiffs allowed to access this information. Plaintiffs also had systems in place to maintain confidentiality with respect to third parties, by, for example, limiting access to its computers. The confidential information was, at all other times, secured on a password protected computer. The confidential information was not externally accessible through the internet. When that password protected computer was not attended, it was secured in a locked building.

80. For the very limited number of additional occasions the confidential, trade secret information was shared, it was shared with a person or entity subject to contractual obligations regarding confidentiality and/or prohibitions on unauthorized use. For the limited number of occasions BearBox shared the confidential, trade secret information with potential customers, it was done only with BearBox management approval.

81. Defendants misappropriated Plaintiffs' trade secrets when they used BearBox technology, without Plaintiffs' authorization, in at least its Smart Response™ software. Defendants have disclosed and/or used the confidential technical information without the express or implied consent of Plaintiffs. The acts of taking Plaintiffs' confidential information, and using it as Defendants' own, and without permission was improper. Further, Defendant McNamara agreed to abide by Plaintiffs' confidentiality terms, only to violate these terms in secret. Defendant McNamara's deception led to Storms continuing to disclose confidential information to Defendants.

82. As a result of Defendants' misappropriation of Plaintiffs' trade secrets, Plaintiffs' have suffered and will suffer damages or other financial harm in an amount to be proven at trial including, but not limited to, damages for actual loss caused by the misappropriation of the trade secrets, financial loss for any unjust enrichment caused by the misappropriation of the trade secrets, and/or damages caused by the misappropriation measured by imposition of liability.

83. The Defendants willfully and maliciously misappropriated the Plaintiffs' trade secrets, and Plaintiffs are therefore entitled to an award of exemplary damages and an award of reasonable attorney's fees under § 51:1434.

COUNT V
CONVERSION BY LANCIUM, MCNAMARA, AND CLINE

~~59-84.~~ Plaintiffs incorporates the above paragraphs by reference.

~~60-85.~~ Austin Storms, in his capacity as founder and President of BearBox, conceived, developed, and reduced to practice ~~the BearBox Technology.~~ BearBox's technology. Plaintiffs own ~~the BearBox Technology, related~~ BearBox's technology, including system designs, documents, data, and know-how, and related intellectual property. Plaintiffs owned this property during all relevant time periods in this suit. ~~Information on the BearBox Technology was provided to Defendants solely for the purposes of evaluation for a potential business relationship and under strict confidentiality obligations.~~

86. Defendants induced Plaintiffs to provide information and documents regarding BearBox's technology to Defendants by misrepresenting that Defendants were interested in investing in or otherwise collaborating with BearBox. Information and documents regarding BearBox's technology were provided to Defendants solely for the purposes of evaluation for a potential business relationship.

Formatted: Font color: Auto, Not Expanded by / Condensed by

87. Without Plaintiffs' consent, Defendants intentionally and willfully assumed dominion and control over the BearBox Technology by claiming BearBox's technology, including system designs, documents, data, and know-how, and improperly used it as to modify their own in Smart Response™ software, and corresponding system designs, to function as reflected in BearBox's system designs, documents, data, and know-how, and subsequently used, sold, licensed, and procured investments related to, and otherwise monetized, that software for substantial profit.

88. Defendants' unlawful exercise of dominion over confidential Bearbox technology, including system designs, documents, data, and know-how not otherwise found to be a trade secret, or having value beyond the '433 patent. Through value of the trade secret(s), has permanently interfered with Plaintiffs' valuable property rights.

~~61-89.~~ Despite providing Defendants with the system designs, documents, data, and know-how that allowed Defendants to modify their wrongful conduct in obtaining the '433 Patent and claiming the BearBox Technology as their own, the Smart Response™ software, and corresponding system designs, Defendants have wrongfully obtained the purported ability to exclude Plaintiffs and others from using the BearBox Technology. This constitutes not compensated or recognized Plaintiffs for the use of BearBox's technology. Defendants' actions constitute an improper and unauthorized and unlawful conversion by Defendants use of Plaintiffs' property.

Formatted: Font color: Black, Condensed by 0.05 pt

~~62-90.~~ As a result of Defendants' wrongful actions, Plaintiffs improper and unauthorized use of Plaintiffs' system designs, documents, data, and know-how to reconstruct and use Bearbox's technology, Plaintiffs have suffered and will continue to suffer imminent and irreparable damages and other financial harms in an amount to be proven at trial. In particular, Plaintiffs have been damaged by losing valuable intellectual property from which As such, Plaintiffs would have derived substantial revenue via licensing and/or selling patented products are entitled to damages resulting from Defendants' improper and unauthorized use.

COUNT ~~IV~~ VI UNJUST ENRICHMENT BY LANCUM, MCNAMARA, AND CLINE

91. Plaintiffs incorporate the above paragraphs by reference.

~~63.1. Plaintiffs incorporate the above paragraphs by reference.~~

~~64-92.~~ Plaintiffs conferred a benefit on Defendants by providing them valuable intellectual property about cryptocurrency mining systems and related confidential information and materials under the boundaries of a potential collaboration between BearBox and Lancium technology, specifically BearBox's technology, including system designs, documents, data, and know-how.

~~65-93.~~ Defendants accepted that cryptocurrency mining intellectual property and, indeed, continuously asked Defendants pressed Storms to provide more information and materials, having recognized the benefit that Defendants received by having access to the BearBox Technology system designs, documents, data, and know-how reflecting BearBox's technology.

94. Defendants accepted induced Plaintiffs to provide these system designs, documents, data, and know-how by misrepresenting that Defendants were interested in investing in or otherwise collaborating with Bearbox. Defendants' actions in obtaining Plaintiffs' system designs, documents, data, and know-how were deceptive.

~~66-95.~~ Defendants retained the BearBox Technology BearBox's technology, including system designs, documents, data, and know-how, and used it to their own advantage, and at Plaintiffs' BearBox's expense.

~~67-96.~~ Defendants have been and continue to be unjustly enriched by profiting from their wrongful conduct. In particular For example, without Plaintiffs' consent, Defendants have unlawfully used Plaintiffs' property by asserting inventorship over the BearBox Technology system designs, documents, data, and know-how to modify their Smart Response™ software to function as BearBox's technology did. As a result, Defendants are deriving an unjust benefit from exploiting Storms's cryptocurrency mining inventions. It would be inequitable for Defendants to retain these benefits under these circumstances. BearBox's property.

97. Despite providing Defendants with the system designs, documents, data, and know-how that allowed Defendants to modify their system designs and Smart Response™ software to function as BearBox's technology did, Defendants have not compensated Plaintiffs for the use of BearBox's technology. It would be inequitable for Defendants to retain these benefits under these circumstances.

Formatted: Font color: Auto

Formatted: Font color: Auto

98. Defendants' unauthorized use of Bearbox's technology has impoverished Plaintiffs by depriving them of the fruits of their labor. In addition to the costs incurred and labor expended in creating Bearbox's technology, Plaintiffs are no longer able to obtain any economic benefit from any use of Bearbox's technology.

99. Defendants' enrichment through the use of Plaintiffs' system designs, documents, data, and know-how to reconstruct and use Bearbox's technology is directly connected to Plaintiffs' impoverishment because Defendants have failed to compensate Plaintiffs for their efforts.

100. There is no justification at law or in contract for Defendants' uncompensated use of Plaintiffs' unique and novel invention.

68.101. Because Plaintiffs have incurred, and continue to incur, detriment in the form of loss of money and property as a result of Defendants' wrongful use of Plaintiffs' intellectual property, including the right to any patent based on their own intellectual property. The intellectual property, including the right to any patents based on Plaintiffs' intellectual property and to any patent documents (including assignment documents), U.S. and foreign, are unique and there is no adequate remedy at law system designs, documents, data, and know-how reflecting BearBox's technology, Plaintiffs are entitled to compensation to the extent Defendants have been enriched or Plaintiffs have been impoverished.

69. The harm to Plaintiffs is continuous, substantial, and irreparable.

COUNT V **NEGLIGENT MISREPRESENTATION BY LANCUM AND MCNAMARA**

70.1. Plaintiffs incorporate the above paragraphs by reference.

71. In connection with the potential work involving cryptocurrency mining systems and related methods, Storms told Defendant McNamara that the cryptocurrency mining systems

and related methods were proprietary to Plaintiffs and not to be used or shared outside of Laneium. Defendant McNamara gave his word that he would abide by this confidentiality. On information and belief, Defendant McNamara agreed to keep the BearBox Technology confidential despite later recklessly incorporating the BearBox Technology into his own patent applications and swearing, as recently as December 4, 2019, that he is an inventor of the BearBox Technology. Storms relied on Defendant McNamara's assurances of confidentiality and continued to share details about the BearBox Technology with Defendants.

72. — If Plaintiffs had known that Defendants would secretly incorporate the BearBox Technology into Defendants' own patent applications to claim them as Defendants' intellectual property, Plaintiffs would not have continued working with and sharing intellectual property with Defendants.

73. — Plaintiffs suffered a pecuniary loss based on this reliance including the loss of potential patent rights, and the costs of Plaintiffs' know-how converted under the guise of a potential business relationship.

JURY DEMAND

74.102. — Under Rule 38(b) of the Federal Rules of Civil Procedure, Plaintiffs respectfully demand a trial by jury on all issues so triable.

PRAYER FOR RELIEF

WHEREFORE, BearBox respectfully requests the following relief:

A. An order that the Director of the United States Patent and Trademark Office correct the inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

B. Alternatively, an order that Defendants sign the requisite documents to correct inventorship of the '433 Patent to name Austin Storms as the sole inventor, or, in the alternative, as a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

C. A declaration that Austin Storms is the sole inventor, or, in the alternative, is a joint inventor to one or both of the individuals currently listed as inventors on the '433 Patent;

D. A preliminary and a permanent injunction enjoining Defendants Lancium, McNamara, and Cline from asserting that McNamara or Cline are inventors of the '433 Patent in violation of the United States federal patent laws;

E. An order that Defendants immediately transfer to Plaintiffs all right, title, and interest in all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of ~~Plaintiffs' intellectual property~~ BearBox's technology;

F. An order for a constructive trust over all information, patent applications, patents, technology, products, and other materials in the possession, custody, or control of Defendants that wrongfully constitute, contain, were based on, and/or derived in whole or in part from the use of ~~Plaintiffs' intellectual property~~ BearBox's technology or any system designs, documents, data, and know-how reflecting BearBox's technology;

G. Financial relief including damages, consequential damages, disgorgement of Defendants' ill-gotten profits, Defendants' unjust enrichment, ~~reasonable royalty damages, lost profits damages~~, reliance damages, and/or all other appropriate financial relief, all in an amount to be determined at trial, with interest;

H. An award of the amount by which Defendants have been unjustly enriched by their actions set forth in this Complaint ~~and their purported ownership of patents covering Plaintiffs' intellectual property;~~

I. A finding that this is an exceptional case warranting imposition of attorney fees against Defendants and an award to Plaintiffs of its reasonable costs and attorney fees incurred in bringing this action pursuant to 35 U.S.C. § 285; ~~and~~

J. An award of exemplary damages and/or an award of reasonable attorney's fees under 18 U.S.C. § 1836(b)(3)(C), (D) and/or under La. R.S. § 51:1434; and

K. An award of such further relief at law or in equity, such as preliminary and/or permanent injunctive relief, as the Court deems just and proper.

ASHBY & GEDDES

/s/ Andrew C. Mayo

Andrew C. Mayo (#5207)
500 Delaware Avenue, 8th Floor
P.O. Box 1150
Wilmington, DE 19899
(302) 654-1888
amayo@ashbygeddes.com

*Attorneys for Plaintiffs
BearBox LLC and Austin Storms*

Of Counsel:

Benjamin T. Horton
John R. Labbe
Raymond R. Ricordati, III
Chelsea M. Murray
MARSHALL, GERSTEIN & BORUN LLP
233 South Wacker Drive
6300 Willis Tower
Chicago, IL 60606-6357
(312) 474-6300

Dated: ~~May 24, 2021~~ February 16, 2022